



WHITE PAPER

Buy behind, sell ahead

Why the mid-market companies furthest behind on AI are often the best ones to buy, and how sponsors close that gap and get paid for it at exit.

October 2026

INTRODUCTION

Most mid-market companies we walk into look much the same. The data sits in four or five systems that disagree with each other, the processes that matter live in a few people's heads, and the AI pilot started last year has quietly stalled. Large companies are moving past this stage: McKinsey's 2026 survey found 40% of companies above \$1 billion of revenue scaling AI agents, software that carries out multi-step work on its own, against 22% of smaller ones.

To a sponsor that gap is the opportunity, because a lower middle market company is priced on its earnings, its size and its sector, and nothing in that price measures how far it is from using AI. Buy a company that is behind, fix the foundations in the first year, and the gain lands on a multiple you've already paid.

This paper draws on what we've built inside PE-backed companies. It covers four moves: buying the gap, integrating add-ons with tools that get reused, owning what a buyer will inspect, and planning for the AI that will exist at exit. Like our June paper on the company brain, it's opinionated; where our view differs from common practice, we say so.

THE GAP IS THE TRADE

The companies furthest behind on AI are often the cheapest to fix, because what holds them back is the foundations.


The pattern we meet most often is a business whose operations work and whose data does not. A storage operator we worked with had valuable customer data trapped in its CRM and point-of-sale systems, with no data team and no plan for either. A multi-brand services company had the same customer under as many as fifteen different IDs. A PE-backed education company wanted AI in its support operation, but its workflows were undocumented and nothing was routed, so there was nothing for AI to build on.

None of these companies needed a better AI model: each needed its data joined, its definitions agreed and its main processes written down, and in each case that took months. The storage operator went on to launch a customer data product that brought in \$1.8 million of new revenue. The education company cut the cost of its support operation from \$3 million a year to \$1 million while customer satisfaction rose from 67% to 89%.

That's the trade: a company like this is bought on today's earnings, and the distance between what it is and what it could be is the sponsor's to close. Underwrite the gain at the multiple you pay, and treat any premium at exit as upside.

EXHIBIT 1

Where the opportunity sits



An illustrative framework: the best targets pair repeatable work with messy data, because the data is the part a new owner can fix.

BUILD THE INTEGRATION ONCE

Every add-on arrives with a different data model, but the platform it joins stays the same, so build the target once and let AI do the mapping.

No two add-ons look alike: one keeps its books in QuickBooks on a chart of accounts a bookkeeper built, the next runs NetSuite, and each names its customers, services and revenue its own way. What stays the same is the platform they join, with one chart of accounts, one customer record and one definition of each number. Done by hand, every deal means mapping a new system into that target from scratch.


That makes the first integration the expensive one, because it's where the target gets defined. Bringing the multi-brand company onto one customer record meant writing down the rules: which records belong to the same customer, which system wins when two disagree, and how a new record is matched when it arrives. That work took five months and cut duplicate records by 62%, and none of it has to be redone for the next brand.

From the second add-on, AI turns the work from building into checking. An agent reads the new company's tables, field names and sample records, and drafts the mapping into the platform's model: this field is the customer ID, that account is residential service revenue. Each proposed match carries a confidence score, a person reviews the uncertain ones, and every approved mapping is saved. By the third add-on the agent is working from precedent, because the same systems and the same quirks keep turning up, and the checks that confirm the totals tie out were written the first time.

For a sponsor, the practical version is to define the platform's target model and build the integration toolkit before the second add-on, then track the cost of integrating each one and expect it to fall. Cheaper integration also makes smaller targets worth buying, and they're usually the cheapest companies in the market.

EXHIBIT 2

Integration cost per add-on, by hand against a reusable toolkit



An illustrative curve: integration cost stays flat when each add-on is mapped by hand, and falls once the target model and the library of approved mappings are reused.

OWN WHAT A BUYER WILL INSPECT

Rent the AI models, which improve every few months, and own the company knowledge they run on.

The code, the data pipelines and the business rules behind a portfolio company's AI should sit in the company's own environment, and none of them should be locked inside a vendor's product. The reason is that the model underneath will keep changing.

Two of the largest users of AI have shown what that looks like. Uber went through its annual AI budget in the first three months of 2026, then changed how it uses AI. Frontier models still do the hard, new work, while repetitive jobs such as code reviews go to open models that cost a fraction as much. A test on real work, run every week, decides which model gets which job, and Uber's cost per AI session has roughly halved while its use keeps growing. Airbnb made the same choice in customer service, where its chief executive has said the newest frontier models are used sparingly in production and Alibaba's open Qwen model handles much of the volume because it's fast and cheap.

In both cases the model was the part that changed. What stays is the company's own knowledge and its record of what a good answer looks like. For a PE fund, we built a governed layer over every source the firm used, from call recordings to email and chat, in ten weeks. In its first week, people across the firm asked it more than a thousand questions. The client owns all of the code, and at handover we gave it a memo explaining each tool we chose and each one we set aside, so whoever runs the system next can see why it was built that way.

Our rule for portfolio companies follows from that: rent the model and the computing power, and own what a buyer will open in diligence. That means contracts and pricing history, processes written down as code, the tests that define a correct result, and the record of who can see what. A company built this way can switch AI vendors without starting over, and can show a buyer exactly how its AI works.

EXHIBIT 3

Rent what gets cheaper; own what a buyer will open

RENT, AND SWAP WHEN SOMETHING BETTER ARRIVES

The AI models, and the computing power they run on

The agent platform that runs the workflows

Off-the-shelf business software, such as a helpdesk or CRM

OWN, BECAUSE A BUYER WILL OPEN IT

Contracts, pricing history and customer records

Processes and workflows, written down as code

Tests that define a correct result

Permissions and audit logs

The test for each layer is whether a buyer would open it in diligence.

PLAN FOR THE AI YOU'LL HAVE AT EXIT

Work AI can't do reliably today may be routine before you sell, so plan for both.


The models improve faster than a hold period runs: researchers who track this find that the length of task AI can complete on its own has doubled roughly every four months since 2023. A value creation plan written against today's models won't capture what the company can do by the time it's sold.

We handle this with a habit from our own builds. Before every build we agree a test with the client: a set of real cases with known right answers, frozen at signing, and a pass mark. Sponsors can plan the same way, with work that passes the company's own test today in the base case. Work that doesn't pass yet goes into a second layer at no value, and is re-tested every quarter on the newest models.

What buyers pay for. Buyers discount AI claims they can't check, so the record matters as much as the result. In one client's system, every query was logged from the first week with the user, the answer and whether it helped. We recommend the same for every AI program in a portfolio, with each saving tied to a number the CFO already reports. That record lets the next buyer check the gain the way they check any other earnings.

EXHIBIT 4

The two-layer plan



An illustrative plan: work moves into the base case only when it passes the company's own test, and the second layer is re-tested every quarter.

WHAT WE EXPECT, AND WHERE TO START

PREDICTION	BY WHEN
At least three mid-market sponsors publicly market a strategy of buying companies that are behind on AI	End of 2027
Asking whether a company can switch AI vendors without a rebuild becomes a standard diligence question	End of 2027
Acceptance tests agreed before the build become the norm for AI work in PE portfolios	End of 2027
Buyers still pay no broad premium for AI used only to work faster	Through 2027
Integration cost per add-on becomes a standard measure for buy-and-build platforms	End of 2028

Five things a deal team can do in the first 100 days:

1. Score each target on how repeatable its work is and how messy its data is.
2. Agree the test before any AI build starts, using real cases from the company.
3. Define the platform's target data model, and build the integration toolkit before the second add-on.
4. Write AI access and vendor-switching terms into software contracts at renewal.
5. Plan in two layers, and re-test the second layer every quarter.

GLOSSARY

TERM	WHAT IT MEANS
Agent	Software that carries out multi-step work on its own, such as reading a contract and entering its terms
Platform and add-on	The first company a sponsor buys in a sector, and the smaller companies it then folds in
Data model	How a company's systems structure and name its information: which fields exist, what they mean and how records link
Acceptance test	A set of real cases with known right answers, agreed before a build; the build passes when the test does
Governed layer	A company's knowledge gathered in one place, with access rules enforced on every read
Value creation plan	The sponsor's plan for growing a company's earnings during the hold
Two-layer plan	A value creation plan that counts only work passing today's test, and re-tests the rest each quarter

NOTES

McKinsey, [The state of AI in 2026](#) (August 2026), on companies scaling AI agents

METR, [Time Horizon 1.1](#) (January 2026), on the length of task AI completes on its own

The Pragmatic Engineer, [The Pulse: tech companies move to open AI models](#) (September 2026), on Uber's AI costs and model choices

Fortune, [Brian Chesky on Airbnb's AI models](#) (October 2025), on Airbnb's customer service agent

ABOUT OUTPOST

Outpost is an AI consulting and implementation firm based out of NYC. We work primarily with private equity and middle-market firms. We build systems like the ones described in this paper and are official partners of Anthropic and OpenAI.

To discuss this paper, contact Imtisal Qadir (imtisal@outpost-us.com) or Talha Shaikh (talha@outpost-us.com). A companion paper, *After the demo*, covers how the people who run portfolio companies get AI from a pilot into production.