THEORY VS REAL USAGE

AI is already in every deal team. The edge is in the routine, not the tool

Almost every fund now uses generative AI somewhere in the deal process. Few have turned it into daily routines that change how many situations they see and how fast they decide.

KEY MESSAGE

Instinctive use is the floor: the edge comes from shared routines that let the same team cover the whole market, more deeply.

3 things to know

1. **Instinctive use is the floor, not the goal.** 90% of US M&A leaders already use or pilot AI [1]; summarising faster is table stakes.
2. **The gain is depth and rigour, not speed alone:** a first look can fall from a day to an hour [3], freeing time for the views that carry decisions.
3. **Routines with owners, templates and checks separate leaders from tinkerers**, and usage logs show which routines are real (see No. 02).

EXHIBIT 1 How M&A teams use AI, % of respondents (Deloitte 2026, n=500 US PE, portfolio-company and corporate leaders)

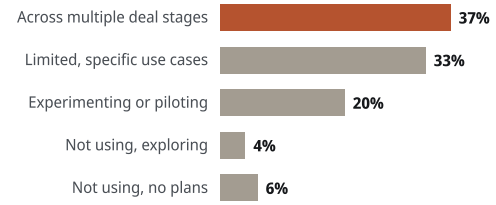

Question covers AI including GenAI. Fieldwork May 2026. Source [1].

90% of US M&A leaders use or pilot AI in M&A, 2026 [1]

37% use it across multiple deal stages [1]

1 day → 1 h screening time per company with GenAI [3]

2/3 of GPs running GenAI pilots, 40%+ in live processes [4]

What deal teams actually do with it

FTI's survey of 208 PE executives finds that all but a small minority use AI in the deal lifecycle. The most common uses are data analysis and summarisation, document scanning, investor profiles, market scanning and investment monitoring [2].

These are reading and writing tasks: market overviews, competitive landscapes, data-room synthesis, first-pass financial checks [5]. They save associate hours. On their own, they rarely change an investment committee outcome.

Theory vs practice

The theory: agents run the deal, from sourcing to IC memo. The practice in most funds: individuals use chat assistants ad hoc, re-check everything, keep prompts in personal notes and leave no record of what was used on which deal.

The gap is not access to tools. Two-thirds of GPs already run pilots and over 40% use GenAI in live processes [4]. The gap is method: shared templates, defined checks, and a clear line between what the machine drafts and what the team decides.

Where the hours go

A first look on a teaser mixes four tasks: reading the document, finding outside information, structuring a view and writing it up. AI compresses the first three; the fourth still needs the person who will defend the view at IC.

The consequence is a shift in the associate's week: less collection, more checking and judgement. Teams that do not adjust roles get faster drafts, not better decisions.

Sources on page 3.

CarveLines © 2026

Five routines that stick

1. Monday pipeline brief: news, filings, hires and funding on the watchlist. **2. First look within 24 hours of a teaser:** business model, market, red flags, questions for management. **3. Data-room triage:** index, missing documents, key contract terms.

4. Expert-call preparation and synthesis. 5. IC memo drafting from the team's own notes, with a source table. Each routine has an owner, a template and a check. Without those three, it remains an experiment.

EXHIBIT 2 Theory vs practice by deal step (CarveLines view; dots = depth of real use in funds today)

STEP	WHAT WORKS TODAY	REAL USE
Sourcing	Watchlists, signals, long lists	●●○
First look	Teaser summary, red flags, questions	●●●
Diligence	Data-room index, contract terms	●●○
Valuation	Model checks, not the model itself	●○○
IC memo	Draft from team notes, source table	●●○
Monitoring	KPI extraction, board briefs	●○○

CarveLines assessment from public surveys [1][2] and field practice.

How to measure it

Not by the number of users or licences. Measure hours per first look, situations screened per quarter, share of IC memos with a sourced fact base, and errors caught at review.

Bain notes that professionals look at about ten deals to find one worth pursuing, often spending a full day on each look [3]. Cutting the first look to an hour frees time for what informs the decision: deeper work, verified facts and better questions on the few that matter.

> "Every PE firm in America has someone on the team experimenting with AI. ChatGPT for market overviews. Perplexity for competitive landscapes. Claude for synthesizing data room documents."
>
> **Rick Claar**, President, Martec Group · 2026 [5]

> "People are nervous about AI. If you don't make it obvious how the technology will improve jobs (rather than eliminate them), frontline employees will refuse to embrace the technology or actively undermine it."
>
> **Bain & Company**, Global Private Equity Report · March 2024 [3]

5 hints: the average fund vs full power

	THE AVERAGE FUND	FULL POWER
1	Summarises teasers on request	**Runs a first-look routine on every teaser within 24 hours**
2	Each person keeps private prompts	**Shared, versioned library of skills per deal step**
3	Accepts outputs that read well	**Requires a source for every number, checked by a person or a second model**
4	Uses AI at the front end only	**Uses it through the cycle: diligence, IC memo, 100-day plan, board briefs**
5	Measures licences and users	**Measures first looks, hours per look, errors caught, decisions changed**

Watch out

Personal accounts on consumer tools. A team member pasting a CIM into a free chatbot is the most common AI risk in a fund today, and the easiest to remove (see No. 08).

Questions for the partners' meeting

1. Which three AI routines does every deal team use each week, and who owns them?
2. What may a team member never paste into an AI tool, and how do we know the rule holds?
3. How many first looks did we do last quarter, and how many could the same team do?
4. Which IC memo of the past year would have changed with a better fact base?

Sources on page 3.
CarveLines © 2026

7 recipes: who, how, what you get

01 24-hour first look

DEAL TEAM

- Drop the teaser into an assistant project loaded with the first-look skill
- Ask for business model, market, red flags and ten questions, each sourced
- Analyst checks every number against its source
- Partner reads two pages and decides: drop, call or deep dive

Output three-page memo
Time 2 hours

02 Monday pipeline brief

ANALYST

- List the watchlist: companies, sectors, people
- Schedule a weekly AI run on news, filings, hires and funding
- One page per sector, three lines per company
- Delivered by 8am every Monday

Output weekly brief
Time half a day to set up

03 Data-room triage

ASSOCIATE

- Load the index and documents into a secure workspace
- Check for missing items against your DD checklist
- Extract change of control, exclusivity and termination terms
- Link every finding to a document and page

Output triage table with references
Time 3 hours

04 Expert-call kit

DEAL TEAM

- AI drafts the call guide from the hypothesis tree
- Record and transcribe the call, with consent
- AI maps answers to hypotheses and flags contradictions
- Update the tree and the next guide

Output updated hypothesis tree
Time 1 hour per call

05 IC memo from notes

DEAL TEAM + PARTNER

- Gather notes, fact base and model outputs in one workspace
- The memo skill drafts in the firm's template
- AI lists every unsupported claim
- The partner edits the argument, not the formatting

Output memo with source table
Time half a day saved

06 Two-model check

ASSOCIATE

- Ask two different models the same high-stake question
- Compare the answers line by line
- Investigate every difference at source
- Record the verified answer in the fact base

Output verified points
Time 30 minutes

07 Monthly routine review

AI LEAD + CHAMPIONS

- Read usage logs for each routine
- Ask each team what was skipped and why
- Fix or retire weak skills
- Publish the updated library

Output updated library
Time 1 hour a month

Monday morning

Ask each deal team for the three AI uses they rely on every week. Turn the top three across teams into written routines with a template, an owner and a check. Review in 30 days: hours per first look and situations screened.

Sources: [1] Deloitte, 2026 Generative AI in M&A Pulse Study (n=500 US PE, portfolio-company and corporate M&A leaders, May 2026); [2] FTI Consulting, AI Radar for Private Equity (n=208), Oct 2024; [3] Bain & Company, Global Private Equity Report 2024, Mar 2024; [4] Pictet 2024 survey, cited by Martec Group, 2026; [5] Martec Group, The impact of AI on commercial due diligence, 2026. Quotes verbatim from the sources cited. Figures as published; survey scopes and dates vary. Vendor-reported figures are labelled; [E] = CarveLines estimate; 'CarveLines view' = our interpretation.

CarveLines © 2026