AI TRANSFORMATION

AI transformation is transformation: full power for everyone, engines inside and out

Funds that treat AI as a tool roll-out get average results. Leaders instrument real usage, raise the bar for every team, fund learning as well as licences, and learn from the frontier.

KEY MESSAGE

Give everyone the full power of AI, then build the engines, inside and out, that turn it into a new way of working.

3 things to know

1. **Full power, for everyone.** No truncated versions: every team member gets frontier models, connectors, automation and, where useful, development tools. AI should raise the bar, not lower it.
2. **Instrument before you invest.** Enterprise usage analytics and logs show who uses AI for what; the gaps show where to tool, train and automate.
3. **Two engines.** Inside: an AI lead, champions with protected time, a budget, incentives, a cadence. Outside: the best start-ups as design partners, specialists, peers.

EXHIBIT 1 **An AI budget has six lines, not one: illustrative split, % (CarveLines estimate [E])**

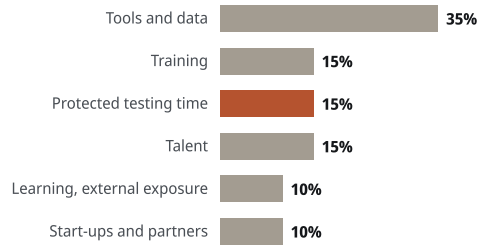

Illustrative allocation for a mid-size fund, CarveLines estimate [E]. Testing time valued at internal cost.

84% of PE firms with a chief AI officer [1]

80% of sponsors that hired have exactly one AI leader [2]

2–3x benefit of redesigning vs optimising, EY client work [1]

~5% of pilots in MIT NANDA's preliminary study delivered rapid revenue gains [3]

Why tool roll-outs stall

Rolling out an assistant and running a training day produces instinctive use: summaries, drafts, quick questions. That is the floor, not the goal.

A preliminary MIT NANDA study found about 95% of the pilots it reviewed delivered no measurable P&L effect [3]; EY warns against "use case syndrome", counting use cases instead of value [1]. Transformation changes how work is done, who does it and what is measured.

Full power, for everyone

Splitting staff into basic and advanced users installs mediocrity at scale. The minimum bar is a power user: everyone, partners included, can build a skill, connect a data source, automate a task and test a result.

Development and data-science tools are not reserved for a tech team. An associate with an AI coding assistant can build in a day what used to need a developer for a month.

Instrument usage

Enterprise AI plans provide admin analytics and audit or compliance exports. Thirty days of anonymised usage data show which tasks repeat, which teams lag, where the same documents or prompts are pasted again and again, and where outputs are reworked.

That is the backlog: tasks to turn into skills, workflows to automate, connectors to add, training to target. Do not stop at what teams do instinctively; compare it with what the best funds and start-ups do.

Sources on page 3.

CarveLines © 2026

The inside engine

One accountable AI lead with a mandate and a budget; a champion in every deal team and in the operating team, with protected time; a shared library of skills and test sets; a monthly session where teams show what they built; and incentives, with AI fluency in reviews and promotions.

Vardis finds 80% of sponsors that hired have exactly one AI leader [2]. One person needs a system around them.

EXHIBIT 2 The transformation engine: components (CarveLines view; dots = effort to build)

COMPONENT	WHAT GOOD LOOKS LIKE	EFFORT
Usage data	Monthly review of usage logs	● ○ ○
Full-power access	Frontier toolset for all staff	● ○ ○
Inside engine	AI lead, champions, cadence	● ● ○
Learning time	Protected hours, tracked output	● ● ○
Outside engine	Start-up design partners, peers	● ● ○
Incentives	AI fluency in reviews	● ● ●

CarveLines view.

The outside engine and the budget

Work with the best start-ups as design partners: early access, influence on the roadmap, and a close view of how AI-native teams work. Run quarterly demo days and peer exchanges. Learn the technology; do not just consume what is available.

Budget six lines, not one (Exhibit 1). Training, learning and protected testing time are investments: track what they produce in recipes adopted and hours redeployed.

> "while companies may have rolled out AI tools, most have not yet productized use cases, redesigned workflows around AI and agentic capabilities, or built the platforms/guardrails needed to run them at scale."
>
> **Alex Singla**, senior partner, McKinsey · Nov 2025 [4]

> "Every single Monday was called 'AI Monday.' You couldn't have customer calls, you couldn't work on budgets, you had to only work on AI projects."
>
> **Eric Vaughan**, CEO, IgniteTech · Aug 2025 [3]

5 hints: the average fund vs full power

	THE AVERAGE FUND	FULL POWER
1	Rolls out a tool and a training day	**Runs a transformation programme with an owner, a budget and a cadence**
2	Gives most staff a basic version	**Gives everyone full power; the minimum bar is power user**
3	Guesses where AI helps	**Reads usage logs monthly and builds the backlog from them**
4	Budgets licences only	**Budgets tools, training, learning, testing time, partners and talent**
5	Buys what vendors offer	**Learns the technology and co-builds with the best start-ups**

Watch out

A one-off training day. Skills fade within weeks without protected practice time, a library to build on and someone who reviews what was built.

Questions for the partners' meeting

1. Does every team member have full-power access, or a truncated version?
2. What do our usage logs say about how AI is really used?
3. Who runs the inside engine, with which mandate and budget?
4. Which three start-ups would we learn most from?

Sources on page 3.

CarveLines © 2026

7 recipes: who, how, what you get

01 Usage-log review

AI LEAD

- Export 30 days of anonymised usage from the enterprise admin console
- Cluster requests by task with AI; count frequency and time
- Spot repeated pastes, reworked outputs and lagging teams
- Rank the ten tasks to turn into skills, workflows or connectors

Output top-10 backlog
Time 2 days

02 Full-power week

ALL STAFF

- Give everyone the full toolset: frontier model, connectors, files, code
- Run a daily 30-minute build challenge on a real task
- Pair each partner with a champion for the week
- Each person ships one personal skill or automation

Output one working skill per person
Time 1 week

03 Champion network

AI LEAD + PARTNERS

- Name a champion in each deal team and in the operating team
- Protect 10% of their time, in writing
- Hold a monthly 45-minute show-and-tell
- Champions own their team's backlog

Output running cadence and backlog
Time 1 week to set up

04 Six-line budget

MANAGING PARTNER

- Split the budget: tools and data, training, learning, testing time, partners, talent
- Value testing time at internal cost
- Set one output KPI per line
- Review spend against output every quarter

Output budget with KPIs
Time 1 day

05 Start-up design partnership

AI LEAD + DEAL TEAM

- Shortlist three start-ups on one core workflow
- Test each on a closed deal
- Agree early access, feedback cadence and data terms with the best
- Send two team members to work alongside them for a day

Output one live partnership per half-year
Time 6 weeks

06 Protected testing time

ALL STAFF

- Block two hours a week per person in calendars
- Pick experiments from the backlog
- Log what was tried, kept and dropped
- Report adoption monthly

Output monthly adoption report
Time ongoing

07 AI fluency in reviews

HR + PARTNERS

- Add a criterion: what you built or automated with AI, and its effect
- Ask for one example with evidence
- Weight it in promotion decisions
- Share the best examples firm-wide

Output updated review grid
Time half a day

Monday morning

Ask your AI lead to export 30 days of anonymised usage data from your enterprise AI plan and present the ten most repeated tasks, the teams that lag and three candidates for automation.

Sources: [1] EY, AI Pulse: US private equity AI insights, 4Q 2025; [2] Vardis Partners, PE AI Leadership Dataset, Aug 2026; [3] MIT NANDA, The GenAI Divide 2025, via Fortune, Aug 2025; [4] McKinsey, The state of AI in 2025, Nov 2025. 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