BUY VS BUILD AND THE FUND'S AI OPERATING MODEL

# Buy the platform, build the context, own the method

Funds are hiring AI leaders fast. The decision that matters is what to own: rarely the model, often the data and the method.

KEY MESSAGE

**Buy platforms and best-of-breed specialists, build the context, and give the AI leader a system and a six-line budget.**

## 3 things to know

1. **Buy platforms and best-of-breed specialists, build context:** vendor solutions succeed about twice as often as internal builds [1].
2. **One AI leader needs a system around them:** champions, budget, cadence; 80% of sponsors that hired have exactly one [3].
3. **The budget has six lines, not one:** tools, training, learning, testing time, partners, talent (see No. 02).

**EXHIBIT 1 Success rate of GenAI deployments by approach, % (MIT NANDA 2025)**

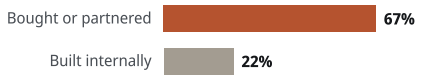


Internal builds succeed 'one-third as often'; 22% is a CarveLines calculation. Source [1].

**84%**
PE firms with a chief AI officer [2]

**80**
fund-level AI hires in 12 months [3]

**350+**
PE, growth and VC firms with an AI leader [3]

**$0.8–1.75M**
AI leader cash pay at mega-funds [3]

## What to buy, what to build

Buy what is generic and fast-moving: models, assistants, document platforms, monitoring software. Configure what is specific: your skills and templates, your connectors to internal data, your evaluation sets.

Build software only when the use case is core, proprietary and stable, as EQT has done with Motherbrain since 2016 (see No. 13). The same logic applies to services: AI-native specialists in research, diligence or value creation now deliver better outcomes than most funds can rebuild.

## Open-source skills libraries

Public libraries of skills and prompts are a fast start: memo templates, Excel checks, research routines. The risks are unknown quality, hidden instructions and drift as models change.

Treat them like third-party code: review, test on past deals, adapt, then version them in your own library.

## Who owns AI in the fund

Three models coexist: a chief AI officer reporting to the managing partner; AI inside the operating-partner team, focused on portfolio companies; or a light central function with a champion in each deal team.

With one leader, the third works best: the leader sets tools, policy and the library; champions turn them into routines.

Sources on page 2.

CarveLines © 2026

## Budgets and lock-in

EY reports AI budgets in PE rising quickly, with a third of firms allocating $50M to $100M [2]. Spend where switching is cheap until value is proven: monthly terms, standard connectors (MCP), exportable prompts and data.

Once a system has learnt your workflows, switching costs become prohibitive [4]. Keep your context portable.

**EXHIBIT 2 Buy, configure or build: a default map (CarveLines view; dots = how much the fund should own)**

| LAYER | DEFAULT | OWN |
|---|---|---|
| **Models** | Buy (API or enterprise tier) | ●○○ |
| **Assistants, platforms** | Buy, run a bake-off | ●○○ |
| **Connectors** | Buy standard (MCP), configure | ●●○ |
| **Skills and templates** | Build and version in-house | ●●● |
| **Test sets** | Build from past deals | ●●● |
| **Core proprietary tools** | Build only if core and stable | ●●○ |

CarveLines view, consistent with [14].

## The talent market

Vardis maps more than 450 AI leaders across private equity and a realistic mid-market pool of about 400 people, with strong candidates receiving two to three sponsor approaches a month [3].

Pay ranges from $275K–$475K total cash at sub-$1bn funds to $0.8–1.75M at mega-funds, often with carry [3]. Smaller funds can share an AI leader with the operating team, or buy the capability as a service.

> "Success requires avoiding the pitfall of "use case syndrome," where volume of use cases becomes the measure of value, as volume of disconnected use cases rarely equates to differentiated enterprise value."
>
> **EY**, AI Pulse, US Private Equity · 4Q 2025 [2]

> "Success in the LLM space isn't about building the most sophisticated system. It's about building the right system for your needs."
>
> **Erik Schluntz and Barry Zhang**, Anthropic, Building effective agents · Dec 2024 [5]

## 5 hints: the average fund vs full power

| | THE AVERAGE FUND | FULL POWER |
|---|---|---|
| 1 | Builds software to look innovative | **Builds only what is core and stable** |
| 2 | Hires one AI lead and stops | **Builds a network of champions with protected time** |
| 3 | Locks into one vendor | **Keeps context portable through open standards** |
| 4 | Budgets licences | **Budgets six lines with output KPIs** |
| 5 | Learns from sales decks | **Learns from start-ups as design partners** |

### Watch out

Hiring an AI leader without a mandate. One person with no budget, no authority over tools and no deal-team time becomes a help desk.

### Questions for the partners' meeting

1. Who owns AI tools, policy and the skills library, with which budget?
2. Which layer are we building that we could buy?
3. How portable is our context if we change vendor?
4. Does each deal team have a named champion?

## Monday morning

Decide in writing who owns AI tools, policy and the skills library, and what budget they control. Share it with the partners and the deal teams this week.

Sources: [1] MIT NANDA, The GenAI Divide: State of AI in Business 2025, via Fortune, Aug 2025; [2] EY, AI Pulse: US private equity AI insights, 4Q 2025; [3] Vardis Partners, PE AI Leadership Dataset, Aug 2026; [4] MIT NANDA, via Technology Magazine, Sep 2025; [5] Anthropic, Building effective agents, Dec 2024. 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