ADVANCED AI FOR PE

Connect the model to your data: MCP, context and the harness decide the result

The same model gives a generic answer or a fund-grade one depending on what it can reach and how its context is managed.

KEY MESSAGE

Reach and context decide the answer: connect the model to your data with open standards and keep its context lean.

3 things to know

1. **MCP is the standard way to connect AI to data:** adopted by OpenAI and Google, 64,935 registry servers [1][3].
2. **Context quality beats context size:** recall falls as context fills [5].
3. **Standardise underneath, specialise on top:** shared connectors and skills, several harnesses.

EXHIBIT 1 **MCP servers listed in the official registry**

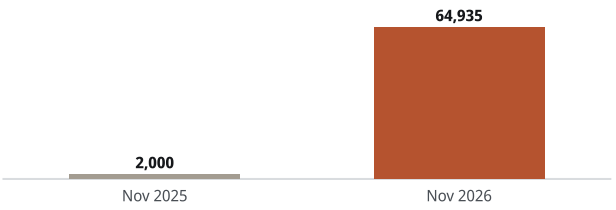

Nov 2025 'close to two thousand'. Source [3], 15 Nov 2026.

64,935 MCP servers in the official registry [3]

656M MCP SDK downloads in 30 days [3]

Mar 2025 OpenAI adopts MCP [1]

Dec 2025 MCP moves to the Linux Foundation [2]

What MCP is

The Model Context Protocol, introduced by Anthropic in November 2024, is an open standard for connecting AI systems to tools and data [2]. A connector written once works with any compliant assistant.

For a fund, it means the CRM, the document store, PitchBook or an internal database can be reached by the assistant with the user's permissions, instead of copy and paste.

APIs and external data

Behind each connector sits an API and a licence. Check three things: entitlements (does the AI respect who may see what), cost (API calls may be billed separately) and audit (are queries logged).

Bloomberg adopted MCP as an organisation-wide standard and cut time-to-production from days to minutes [3]. The gain comes from standardisation, not from the model.

Context management

Models degrade as their context fills: recall falls as tokens accumulate, which Anthropic calls context rot [5]. Good systems give the model only what the task needs.

They use sub-agents that explore widely and return short summaries [5], and keep firm knowledge in skills that load on demand. More context is not better context.

Sources on page 2.

CarveLines © 2026

Choosing a harness

The harness is the software around the model: Claude, ChatGPT, Copilot, a vendor platform or your own. Judge it on support for open standards, quality of tool use, control over data and logs, and the ability to load your skills.

Avoid locking firm knowledge into a format only one harness reads.

EXHIBIT 2 Four layers of a fund's AI stack (CarveLines view; dots = how much to standardise)

LAYER	WHAT IT HOLDS	STANDARDISE
Harness	Claude, ChatGPT, Copilot, vendor	●○○
Skills	Methods, templates, checks	●●●
Connectors	MCP to CRM, data, documents	●●●
Data	Licensed, internal, deal files	●●○
Controls	Permissions, logs, retention	●●●

CarveLines view, consistent with [25].

Interoperability and specialisation

No single harness is best at everything. A fund may use one for research, one for documents and one inside Excel, all sharing the same connectors and skills.

Open standards make this possible: MCP for data and tools, plain-text skills for methods. Specialise the front end; standardise what sits underneath.

> "LLMs are most useful when connecting to the data you already have and software you already use."
>
> **Mike Krieger**, Chief Product Officer, Anthropic · Mar 2025 [1]

> "People love MCP and we are excited to add support across our products."
>
> **Sam Altman**, CEO, OpenAI · Mar 2025 [1]

5 hints: the average fund vs full power

	THE AVERAGE FUND	FULL POWER
1	Copy-pastes data into chat	**Connects sources with permissions and logs**
2	Loads everything into context	**Loads what the task needs; uses sub-agents**
3	Uses one harness for everything	**Uses the best harness per job on shared connectors**
4	Keeps skills in a vendor format	**Keeps them in portable plain text**
5	Lets anyone add connectors	**Approves connectors with least privilege**

Watch out

Connecting everything. Each connector widens what the model can see and what can leak. Connect by use case, with least privilege.

Questions for the partners' meeting

1. Which of our data sources can our assistant reach today?
2. Do our connectors respect the same permissions as our users?
3. Could we change harness without rewriting our skills?
4. Who approves a new connector?

Monday morning

Ask your AI lead which data sources your assistant can reach today, under which permissions, and which three should come next.

Sources: [1] TechCrunch, OpenAI adopts rival Anthropic's standard for connecting AI models to data, 26 Mar 2025; [2] Wikipedia, Model Context Protocol, accessed Nov 2026; [3] Deepak Gupta, MCP enterprise adoption guide, 15 Nov 2026; [4] PitchBook, press release on LLM partnerships, 22 Jul 2025; [5] Anthropic, Effective context engineering for AI agents, Sep 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