From AI Pilots to Performance

Scaling AI Across the Private Equity Investment Lifecycle

HCPEA | October 7, 2026

Agenda

01	**Opening**	Where PE AI stands today and the Crawl, Walk, Run model
02	**Crawl**	Individual fluency: tools, connectors, add-ins, starter tasks
03	**Walk**	Team skills, review gates, portco prototypes, live demo
04	**Run**	Context layer, managed agents, and the team behind them
05	**Getting started**	A 90-day plan from any starting point
06	**Q&A**	Questions and discussion

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

OUR TEAM



Kate Festle
Head of Healthcare M&A
kfestle@westmonroe.com



Keith Campbell
M&A Practice Lead
kcampbell@westmonroe.com



Joe Vattimo
M&A Advisor – Applied AI
jvattimo@westmonroe.com

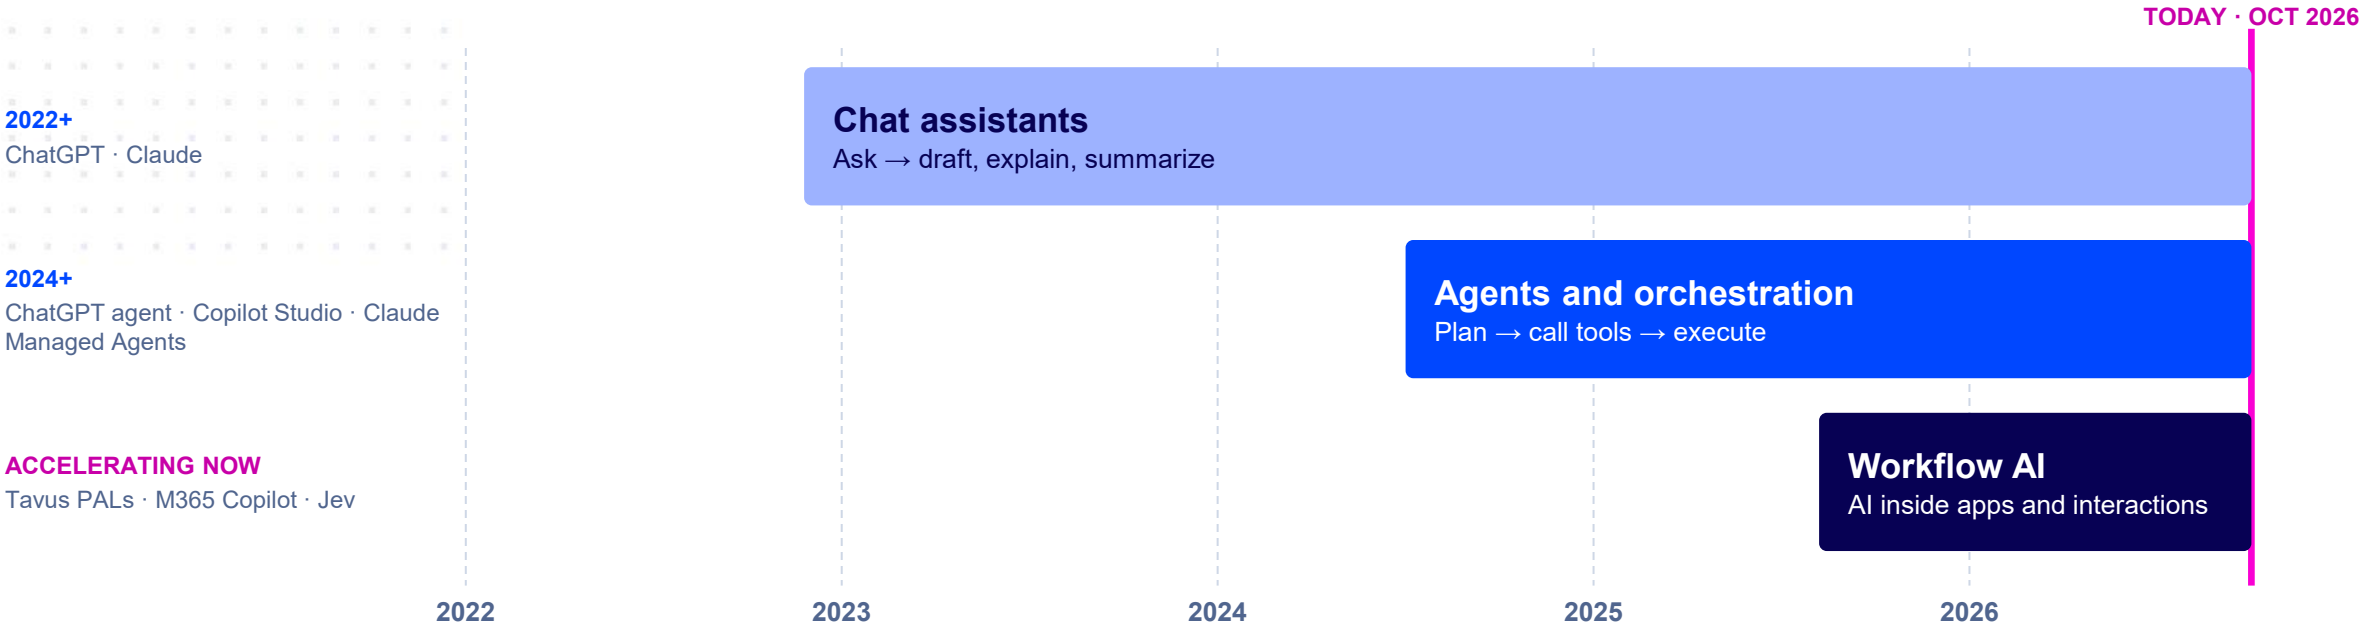


Noah Alex
Engineer – Applied AI
nalex@westmonroe.com

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

AI DIRECTION OF TRAVEL

AI continuously improves our daily work, with each wave adding to the last



POTENTIAL PE USE CASE · PROTOTYPE

Tavus IC sparring partner

A live video agent, grounded in approved deal materials, challenges assumptions before IC. Configure one devil's advocate, or pair it with an evangelist.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

DEFINITIONS

LLM, agent, harness – each are a vital part of how we use AI

HARNESS

Everything around it

Keeps the agent useful and safe.

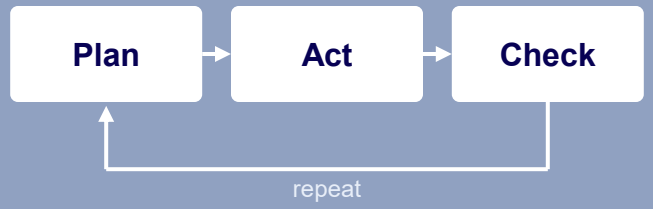
- Instructions
- Skills
- Connectors
- Permissions
- Tests
- Logs

e.g. Claude Code, managed agents

AGENT

The model in a loop

Works until the task is done (and calls tools).

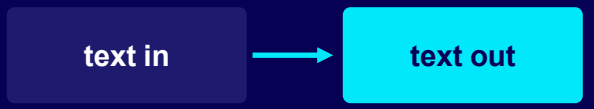


e.g. Cowork, ChatGPT agent

LLM

The model

Reads and writes text but never acts.



e.g. Claude, GPT

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

PE AI GOVERNANCE

How a firm enables AI across the firm (more autonomous vs. more controlled) creates ripple effects that impact maturity and speed

← MORE AUTONOMY

MORE CONTROL →

OPEN ACCESS

- \+ Faster innovation
- \+ Learning compounds
- − Data exposure
- − Tool sprawl and cost

THE SWEET SPOT

Risk-tiered guardrails

LOW RISK
Open: research, drafting

MEDIUM RISK
Approved skills only

HIGH RISK
Gated: deal data, MNPI

CONTROLLED ACCESS

- \+ Safer data boundaries
- \+ Auditable, consistent
- − Slower experimentation
- − Missed frontline ideas

WEST MONROE OBSERVATION | Smaller funds = less control and Broader, multi-strategy AUM funds = more control

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

PE FUND ROLES

Thinking about AI tooling as jobs to be done helps to shape how it should be governed and utilized across the investment lifecycle

ROLE	EXAMPLE BASIC RESPONSIBILITIES	EXAMPLE AI TOOLS
Analyst / Associate	• Build LBO / returns cases; prep IC exhibits • Screen targets; size markets; compile comps	• PitchBook / CapIQ: targets, comps, deal data • AlphaSense: research • Claude: memo synthesis, CIM analysis, Excel / PowerPoint
Vice President	• Run diligence, advisor workstreams, VDR Q&A • Track IC actions, approvals, dates, process risks	• DealCloud: stages, owners, diligence tasks • Claude / ChatGPT: VDR analysis, IC materials development • Krisp / Notion: call actions / recording, agent building
Principal	• Lead thesis, valuation, structuring, negotiation • Turn underwriting into 100-day plan / board agenda	• AlphaSense / Bridgetown: thesis / diligence research • PitchBook / CapIQ: precedents, market/funding research • Claude / ChatGPT: scenario review, 100-day plan build
Partner	• Approve deals / exits; set capital allocation • Review NAV, exposure, fund / asset returns	• Chronograph — fund / asset financials, NAV • DealCloud (pipeline); Notion (decision log)
Operating Partner	• Review revenue, margin, cash, working-capital KPIs • Set initiative owners, milestones, value capture • Value creation plan build and execution	• Chronograph: KPI collection / portfolio tracking • Notion / Krisp: ops reviews, owners, actions • Frontier Models: skills, research, VC planning

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

FRAMEWORK

Each level reflects changes in how AI is used, governed, and encouraged

	CRAWL Individual	WALK Team	RUN Fund
Who uses it	Individuals	Deal and ops teams	Whole fund + portcos
Tools	Chat + add-ins	Skills + agentic workspaces	Managed agents + firm apps
Context	Attached files	Shared connectors	Governed context layer
People	Self-serve + sponsor	Champions + first FDI	Standing FDI pod
Governance	Acceptable-use policy	Review gates + owners	Access, evals, audit
Proof of value	Hours saved	Cycle time	EBITDA traced to AI-enabled work

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

01

Crawl

Individuals get fluent: chat, connectors, and add-ins applied to work they already own

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

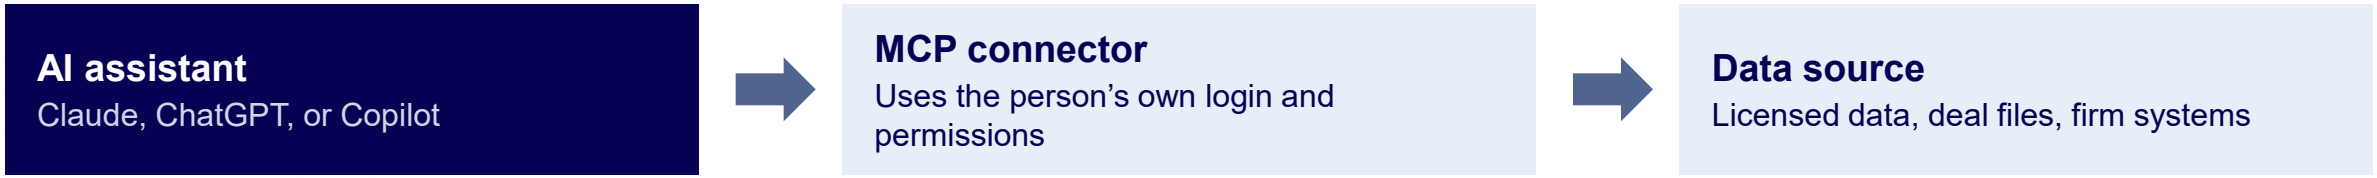
Five tool types cover most deal team work today; match each one to the job

TOOL TYPE	HOW IT WORKS	USE IT FOR
Chat assistants ChatGPT, Claude	You attach files or name sources; the model reads them and answers within one conversation/session.	Research, ideation, document review, and first drafts
In-app agents Microsoft 365 Copilot, Claude for Excel and PowerPoint	Runs inside the open file and reads or edits the cells, slides, or text you are reviewing.	Model checks, IC slide edits, memo redlines
Agentic workspaces Claude, Codex, Other	Plans a multi-step task, works across files and connected apps, and returns output for review.	File inventories, workbook builds, recurring reports, application builds
Call recording and notes Notion AI Meeting Notes, Krisp, Granola, Otter	Transcribes the call, then drafts a summary, decisions, and action items into a searchable record.	Management meetings, expert calls, portfolio check-ins
Specialist platforms Grata, Hebbia, licensed financial data	Purpose-built search and analysis over curated or licensed datasets.	Sourcing lists, large document sets, market evidence

*Before recording, confirm participant consent and check expert-network/NDA terms.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Connectors (MCP) let the assistant query licensed and firm data directly instead of pasted files



CATEGORY	EXAMPLE CONNECTORS AVAILABLE TODAY	EXAMPLE JOB
Private markets and deals	PitchBook, S&P Capital IQ, Chronograph, SS&C Intralinks	Pull deal history and comps into a screening memo
Market data and fundamentals	FactSet, LSEG, Morningstar, Moody's, MSCI, Daloopa	Refresh public comps and credit data in the model
Research and experts	Third Bridge, Guidepoint, GLG, IBISWorld, Aiera	Summarize expert views on a thesis question
Firm systems	SharePoint, Box, Egnyte, Snowflake, Databricks; CRM through a custom MCP server	Answer from the deal folder and pipeline records

Use connectors well: start read-only, keep each user on their own login, and ask for citations. Limit connectors in a single skill or create separate skills with explicit instructions for how each should be used.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Add-ins put the assistant inside the model and the IC deck, where analysts already work



Excel

Claude for Excel
M365 Copilot in Excel

TRY FIRST

- Build an LBO from a template
- Pull comps through a data connector
- Create business case analysis for value creation plan



PowerPoint

Claude for PowerPoint
M365 Copilot in PowerPoint

TRY FIRST

- Turn approved analysis into IC slides on the firm template
- Rewrite titles as findings
- Chart a table in the deck's style



Word and Outlook

Claude in M365
Copilot in Word and Outlook

TRY FIRST

- Redline a memo against the IC format
- Draft LP and banker follow-ups for review
- Turn a long email thread into actions

General Advice

- **Consider the plug-ins as a broader part of your workflow. They greatly improve slide building / excel modeling vs. chat**
- **Building context in a chat or with a skill and then outputting to plug-ins will yield better results**
- **Push them to the limit – opens up new avenues for data analysis (e.g., customer cubes in diligence) that historically were very time consuming**

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

WORKFLOW | DILIGENCE TO IC

A customer call can become a source-linked IC insight; while AI analysis may not be perfect, it will greatly speed up cycle times and analysis

01

Capture + transcribe

TOOLS | Krisp / Notion

Record a customer interview. Keep the transcript with speaker names and timestamps. Use Notion agents / scheduled agents to parse and organize information

02

File deal evidence

TOOLS | SharePoint / DealCloud

Save transcript beside the CIM and KPI workbook; log follow-up against the deal OR create an automated workflow via agent to save

03

Extract in context

TOOLS | Claude / GPT project + skill

Use the existing context layer to ask questions and structure what can be put into plugins (e.g., Excel, PowerPoint)

04

Build IC outputs & Review

TOOLS | Excel / PowerPoint

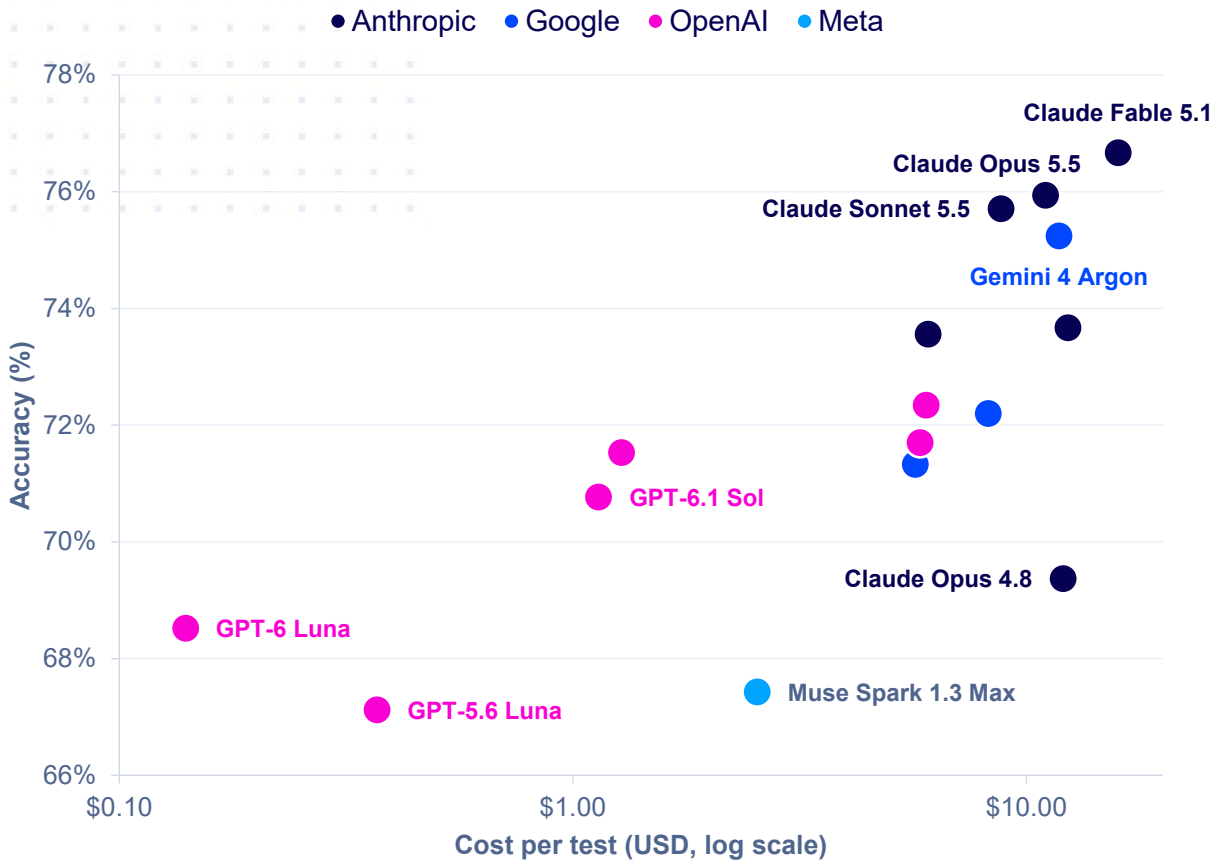
Copy and paste prompts, with context, or drive from Claude / Codex to make file updates. Create schedule tasks or agents to automate

Key Risks & Training to Execute:

1. As AI builds become more and more common, keeping an eye on junior team members work is critical. Look for standard AI written items such as uncommon symbols in writing, significant semi-colons, and unclear writing / rationale. Also look for oddly configured shapes with lots of individual components
2. Take time to train junior resources on what is correct and what isn't. While the work can move faster, sloppy work done quicker is a higher risk for investment decisions
3. Build best practice guidelines, templates, and outputs, that can be used in context to minimize AI errors

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Top models are within 1.5 points on Excel modeling; cost varies 114×



At this point, *all* frontier models are *extremely* capable. Use them to enable every aspect of your current workflows.

PRICE POINTS

- Accuracy is close at the top; cost per test is not: $0.14 to $15.93
- Mid-tier models like Claude Sonnet 5.5 land within a point of the leader
- Low-cost tiers like GPT-6.1 Sol score about 6 points lower at 1/14 the cost

PE TAKEAWAY

- Route by tier: frontier for judgment, cheaper models for volume
- Test on your own CIM, Excel, and IC-deck tasks
- Treat benchmarks as directional

Source: Vals AI, Excel Modeling Benchmark, overall task. Third-party benchmark; cost per test reflects list API prices.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Four analyst tasks to pilot this week

TASK	STAGE	TOOL	REVIEWER CHECK
Build a target long list from a thesis	Sourcing	Chat with PitchBook connector	Criteria met; deduplicated against CRM
Summarize a CIM into a one-page screening memo and feed to a one-page model	Diligence	Connected Chat Project + Plugins	Page reference on every figure; assumptions QA'd
Turn expert call transcripts into a themes grid	Diligence	Call notes tool with chat	Quotes trace to the transcript
Trace and document an LBO model's formulas	Underwriting	Claude for Excel	Reviewer reruns one scenario

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Five crawl-stage pitfalls and how to address them

PITFALL	WHAT IT LOOKS LIKE	FIX
Firm data in consumer tools	Deal files pasted into personal accounts	Use an enterprise-approved tenant; follow firm data and retention controls
Context issues	Thin deal context, chat drift across files or versions, and unsourced numbers can produce generic or incorrect output	Start a fresh task with approved deal files and thesis; cite claims and reconcile figures before IC
Prompts nobody reuses	Each analyst rewrites the same prompt	Share tested prompts and turn repeatable tasks into maintained skills
Measuring logins, not outcomes	High usage, no change in cycle time	Track hours, cycle time, quality, and rework on named tasks
Recording call blocks	Management or expert calls are not transcribed under firm policy	Use transcription only when consent, law, and firm policy permit; otherwise use approved notes

Top barriers PE firms report: lack of expertise (49%), data privacy (43%), and model accuracy (38%). Source: S&P Global Market Intelligence, 2026 PE and VC Outlook survey (February 2026).

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

02

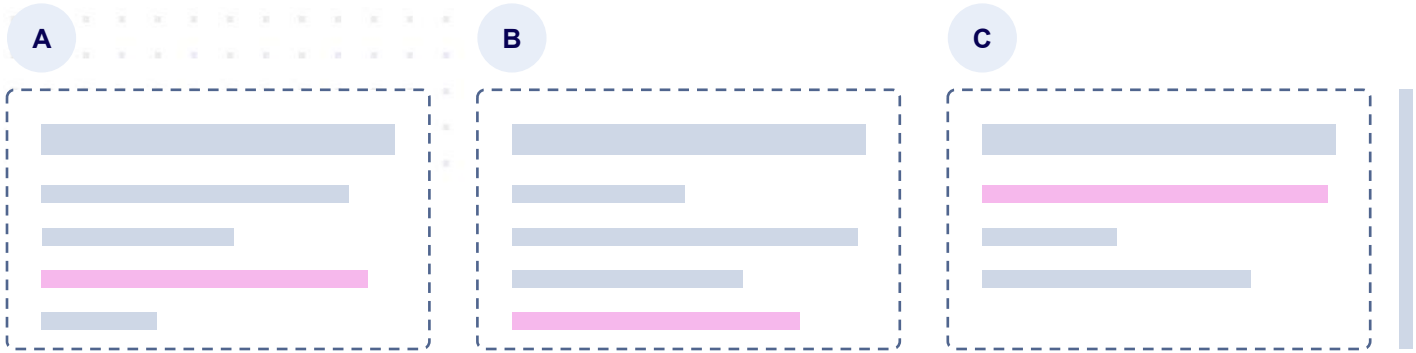
Walk

Teams share what works: skills, agentic workspaces, and review gates

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

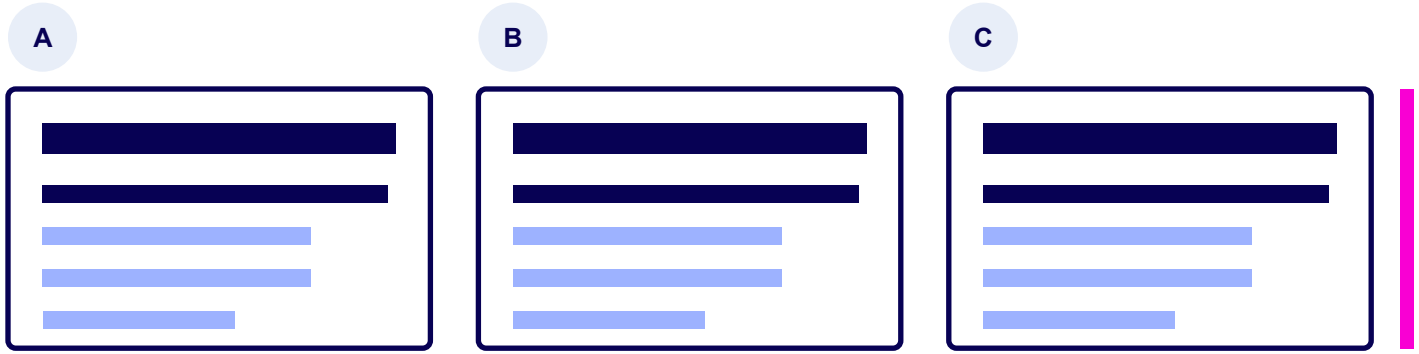
Skills are reusable, sharable files that can create standardization and give others a running start on workflows

WITHOUT A SKILL



Different sections, different numbers, different definitions. Harder to review and verify.

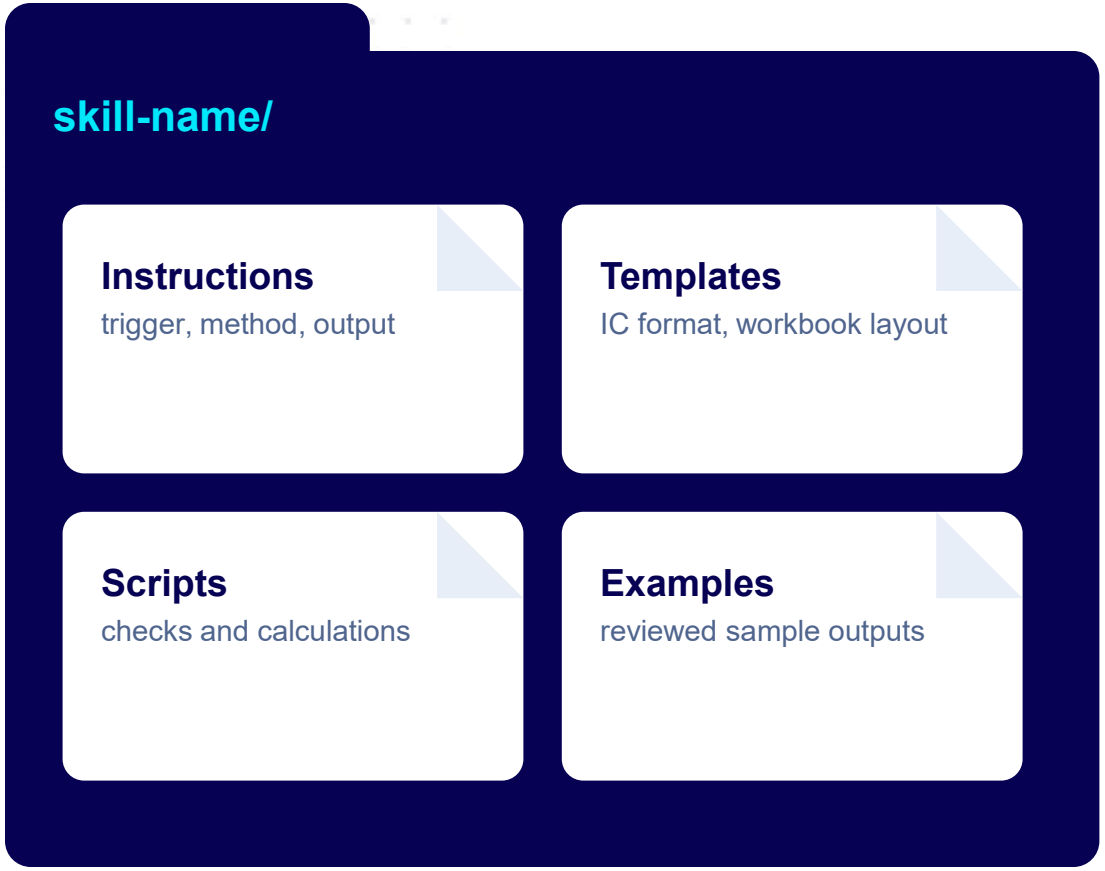
WITH A SKILL

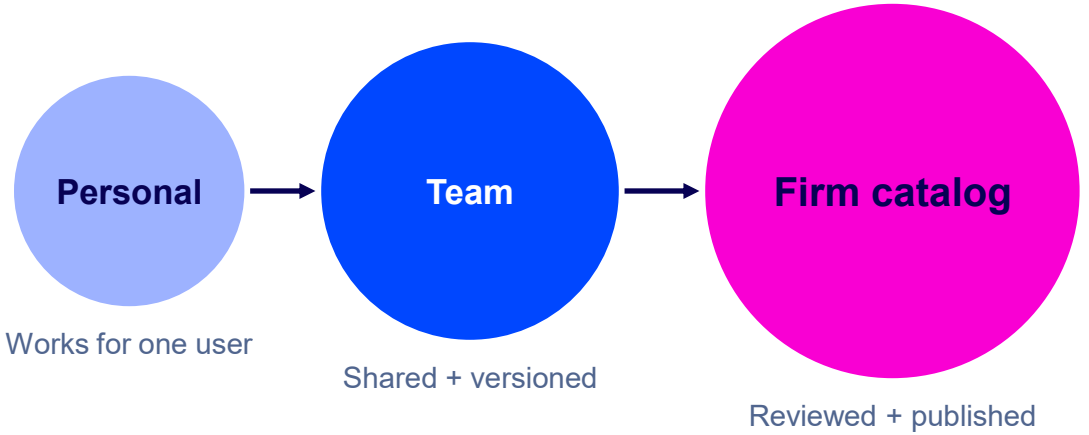


Same sections, same checks, sources cited every time. Easier to review and verify.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

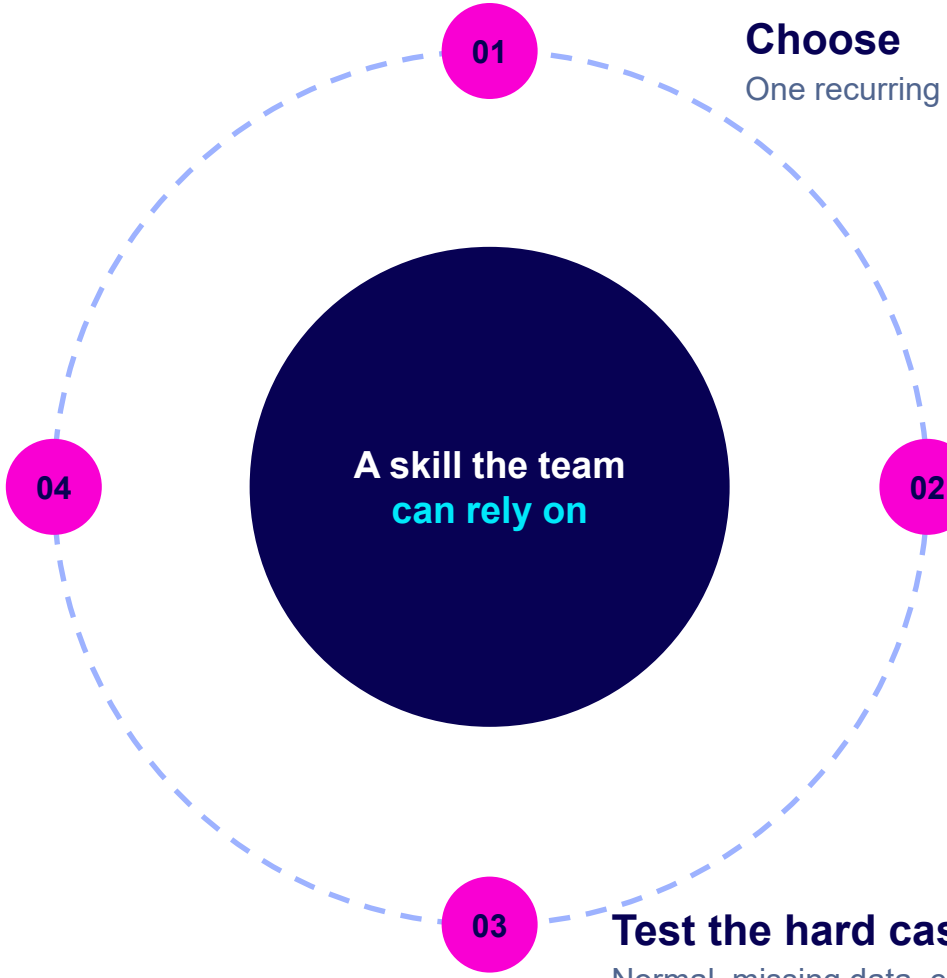
A skill turns a prompt that works into a method the whole firm can use, which can be scaled up and standardized





© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

A reusable skill is built in four steps and tested on the cases with known outputs



01 Choose
One recurring task, like an IC risk register

02 Package
Trigger, inputs, method, template, checks

03 Test the hard cases and cases with a verifiable answer
Normal, missing data, conflicting sources

04 Own and recheck
Named owner and version; retest when anything changes

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

A skill reaches the firm store only after a named reviewer signs off, and the store is governed like code

01 Build

Owner turns a prompt that works into a SKILL.md.

Author: deal team member

02 Sign off

Reviewer runs the test cases, reads the instructions and scripts, and approves.

Reviewer: deal lead, FDI

03 Publish

Approved version moves to the firm repository with an owner and data scope.

Publisher: AI or IT admin

04 Maintain

Usage and errors are tracked; retest when the model, template, or source format changes.

Owner: skill author

Two skill type examples

Skill	What it does	Sign-off check
Extraction /cim-extract	Pulls revenue, EBITDA, and customer data from a CIM into the KPI template	Every figure ties to a cited page
Analytics /ic-variance	Compares actuals to plan and drafts commentary for the IC memo	Reviewer approves each claim before use

Rules for a store of AI-written skills

- **A named human owns every skill**, even when AI drafted it
- **No publish** without passing tests and written sign-off
- **Scripts and connector access** reviewed like code
- **Every version logged**, with one-click rollback
- **Retire skills** unused for 90 days

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Human review scales with risk, from spot-checks on internal drafts to named approval for anything external

TIER 1

Internal and reversible

Call summaries, file indexes, first-draft research

REVIEW

Spot-check a sample each week

TIER 2

Feeds an investment decision

IC memo analysis, model changes, KPI commentary

REVIEW

Reviewer ties numbers to sources and signs off

TIER 3

External or binding

LP letters, valuation marks, bids, board materials

REVIEW

Named approver; AI drafts but never sends or submits

HOW TEAMS BUILD TRUST IN THE OUTPUT

Citations
Every claim links to a page, cell, or record

Exception lists
The agent flags conflicts instead of silently fixing them

Test cases
Past deals with known answers check each new skill version

Audit trail
A log of who reviewed and approved each output

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

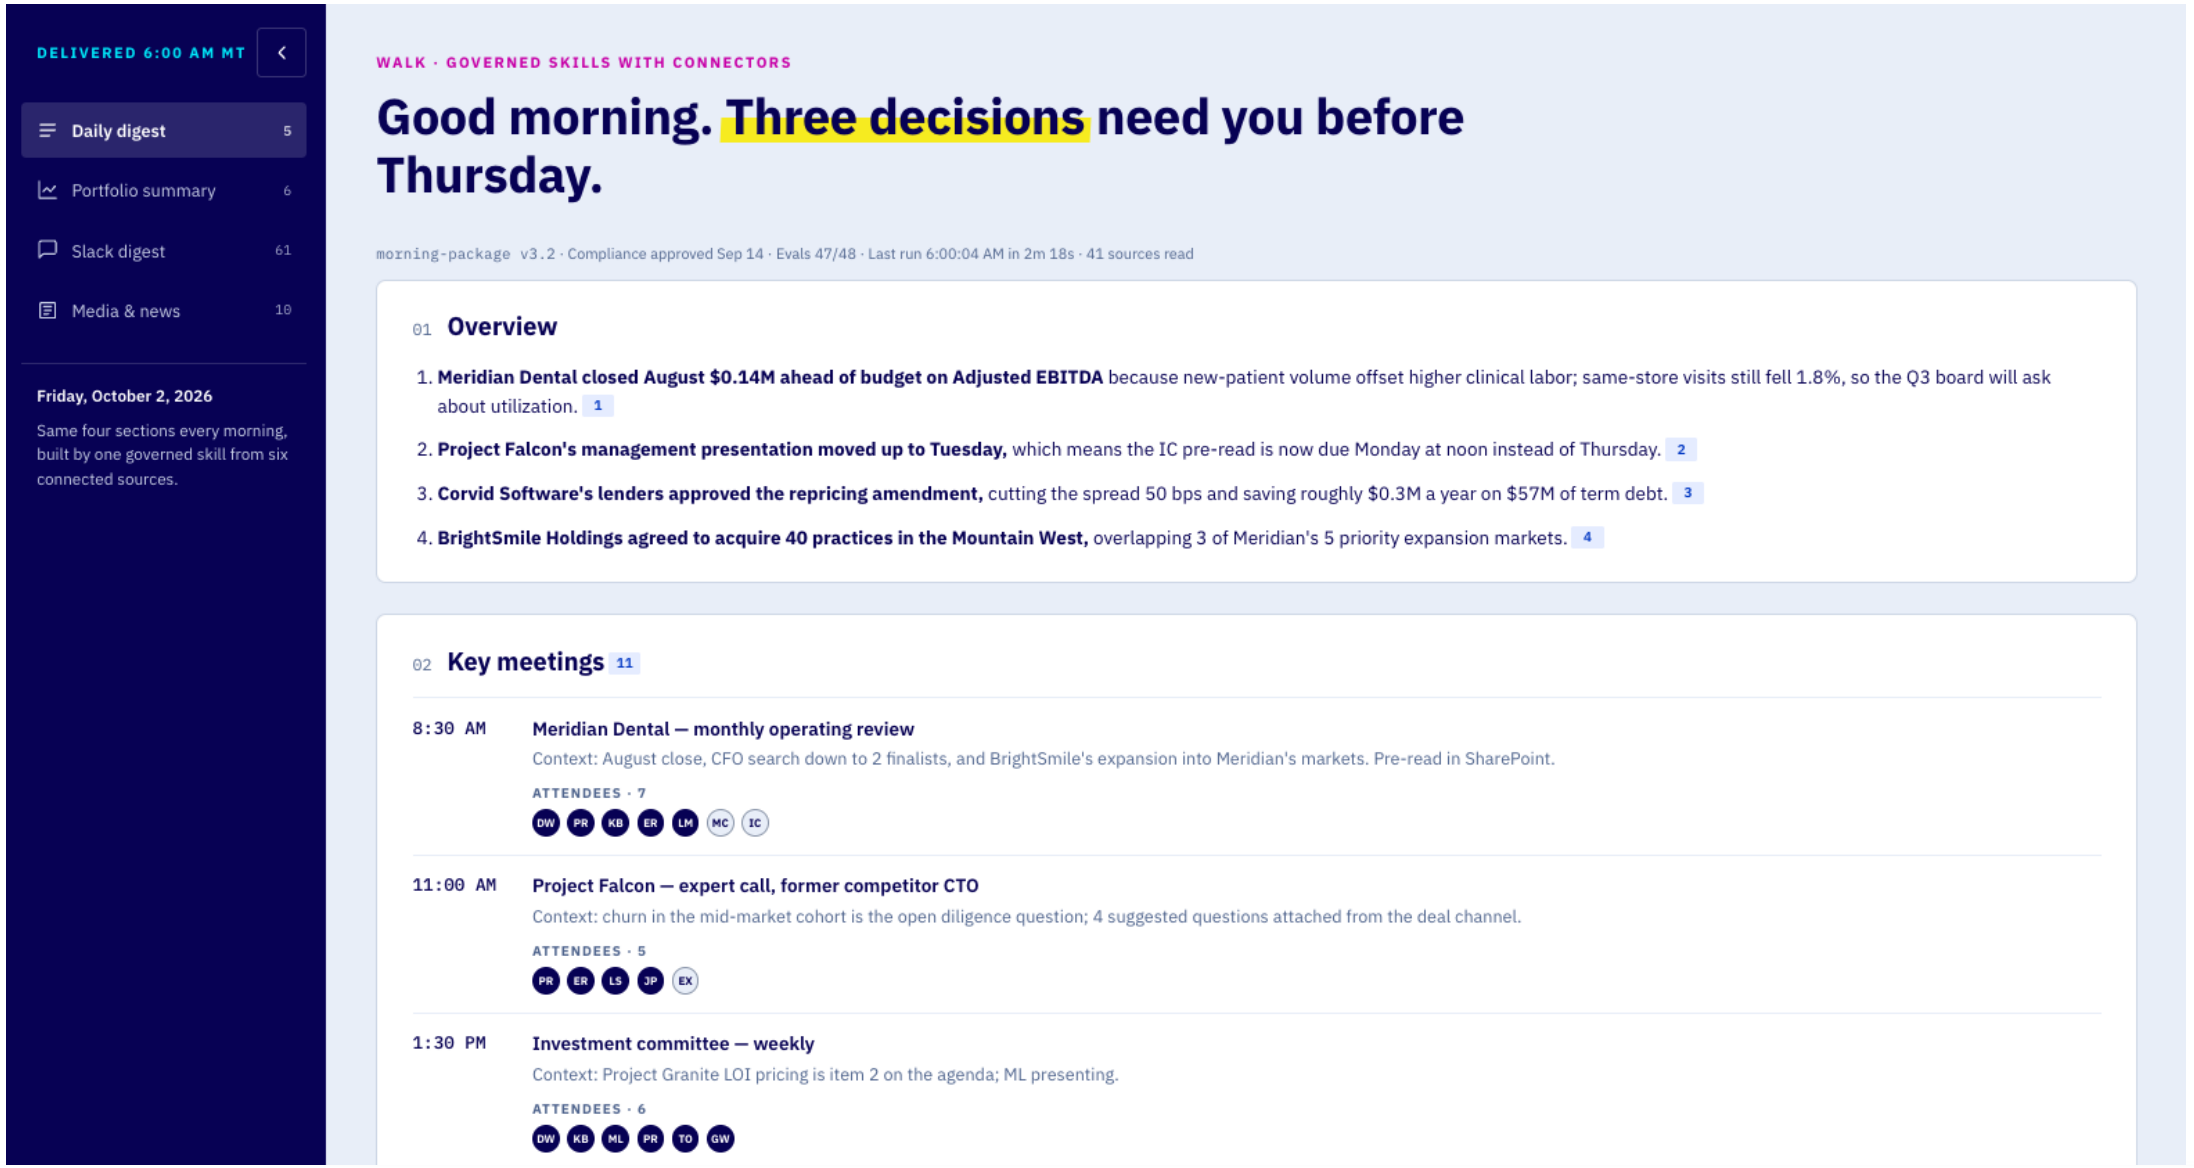
Walk needs a small dedicated team: AI champions in each group and the first Forward Deployed Investor

ROLE	TIME COMMITMENT		WHAT THEY DO
Executive sponsor Partner or COO		**Part time**	Sets priorities, approves spend, removes blockers
AI champions One associate or VP per team		**10–20%**	Collect use cases, maintain team skills, coach peers
Forward Deployed Investor Hire, or convert a technical investor		**Full time**	**Turns recurring deal work into skills and agents; builds portco prototypes**
IT and compliance partner Existing IT, CISO, or CCO		**Part time**	Provisions tools and connectors, reviews permissions, owns the policy
Data engineer (optional) Contractor or shared resource		**As needed**	Connects SharePoint, CRM, and disparate source data

1–2 **dedicated people plus champion time, largely independent of fund size.**

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

"Walk" Demo

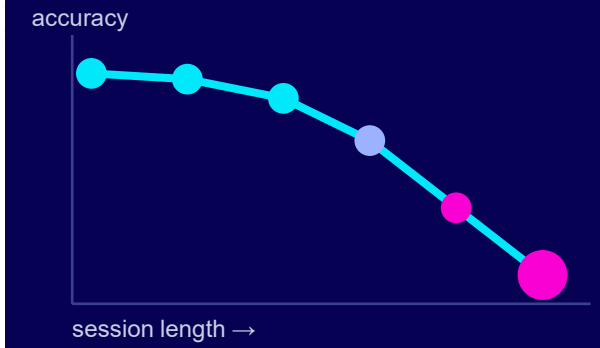


© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Output quality decays quietly when sessions run long or sources conflict

CONTEXT ROT[1]

Accuracy drops as a session fills up



LOOKS LIKE

- Project A details in the Project B memo
- Cites last month's model
- Forgets the IC format

FIXES

- One task per session (/clear or restart with a summary)
- A curated context pack
- Skills load only what a task needs (e.g., DealCloud only)
- Use the correct/larger model

DATA QUALITY

Agents repeat whatever the source says



LOOKS LIKE

- Two CIM versions in one folder
- One company counted twice
- Billed vs. scheduled visits

FIXES

- One approved, dated version per folder
- A definitions file per deal, with a repeatable folder structure
- An owner who refreshes each source on a schedule

KEEP THREE LAYERS OF CONTEXT SEPARATE

Firm criteria, IC format

Deal sources, definitions

Working decisions, open issues

1. Chroma Research, "Context Rot: How Increasing Input Tokens Impacts LLM Performance," 2025.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

03

Run

The fund builds the foundation: a governed context layer, managed agents, and a standing pod

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Run requires a governed context layer; without it, agents stay confined to individual sessions and files

Experiences	Chat, Excel and PowerPoint, Cowork, Claude Code, scheduled agents, and firm apps
Skills and agents	Firm catalog of versioned skills; agents with named owners and test cases
Access (MCP)	Firm MCP servers expose the layer with each user's permissions and log every call
Context layer	Context graph of companies, deals, people, and decisions; document index; metric definitions
Data platform	Snowflake, Databricks, or Fabric: ingestion, entity resolution, quality checks
Source systems	CRM, VDRs and SharePoint, fund admin and portfolio monitoring, portco ERP and EHR feeds, call transcripts, licensed data

Security and governance

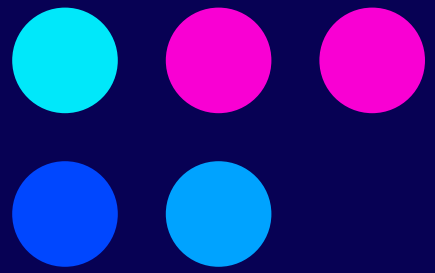
- Single sign-on
- Role-based access
- Fund and portco separation
- Audit logs
- Retention and deletion
- Cost budgets

Build from the bottom up; each layer can start small!

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

A few people build the platform; everyone else builds on it

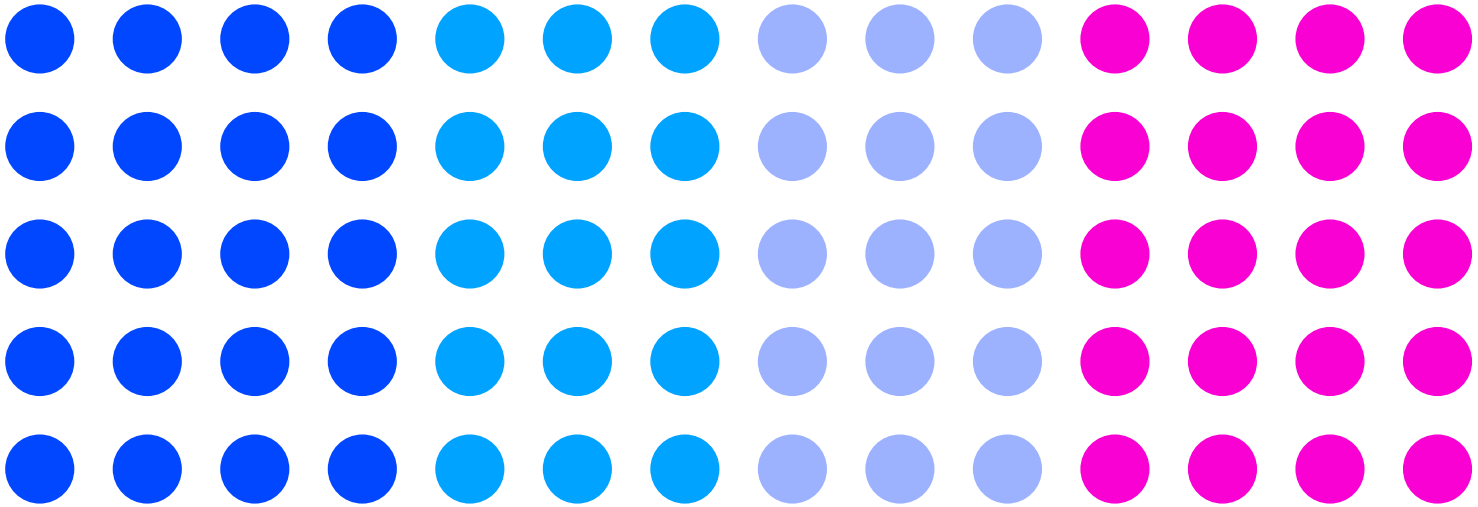
THE POD (FDIs + Data Resources)



2–5

dedicated people at a mid-market fund, scaling to 10 – 15 for a mega fund

builds the platform



Everyone else: builds and owns their own skills and agents (illustrative)

Deal teams | Operating team | IR and finance | Portco teams

Context layer
one governed source of firm data

Catalog and guardrails
reviewed skills, evals, audit

Coaching
FDIs embed with teams to scope work

West Monroe estimate; team size depends on strategy count, portfolio size, and how much of the data platform already exists.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

RAG finds passages; a context graph connects people, deals, and decisions

RAG pipeline

Unstructured document passages, indexed by semantic meaning



“What did the CIM say about payer mix?”

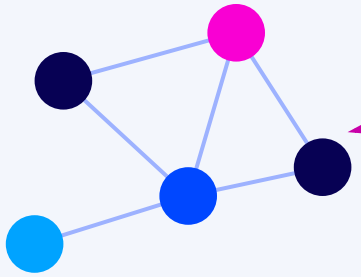
Pro: Large document sets

Con: Not great for certain tasks, such as counting and comparing across deals

Build Time: Weeks

Context graph

Entities, relationships, and decisions over time



“Which bankers sent us healthcare deals, and why did we pass?”

Pro: Questions across deals, people, and time (think ”relationships”)

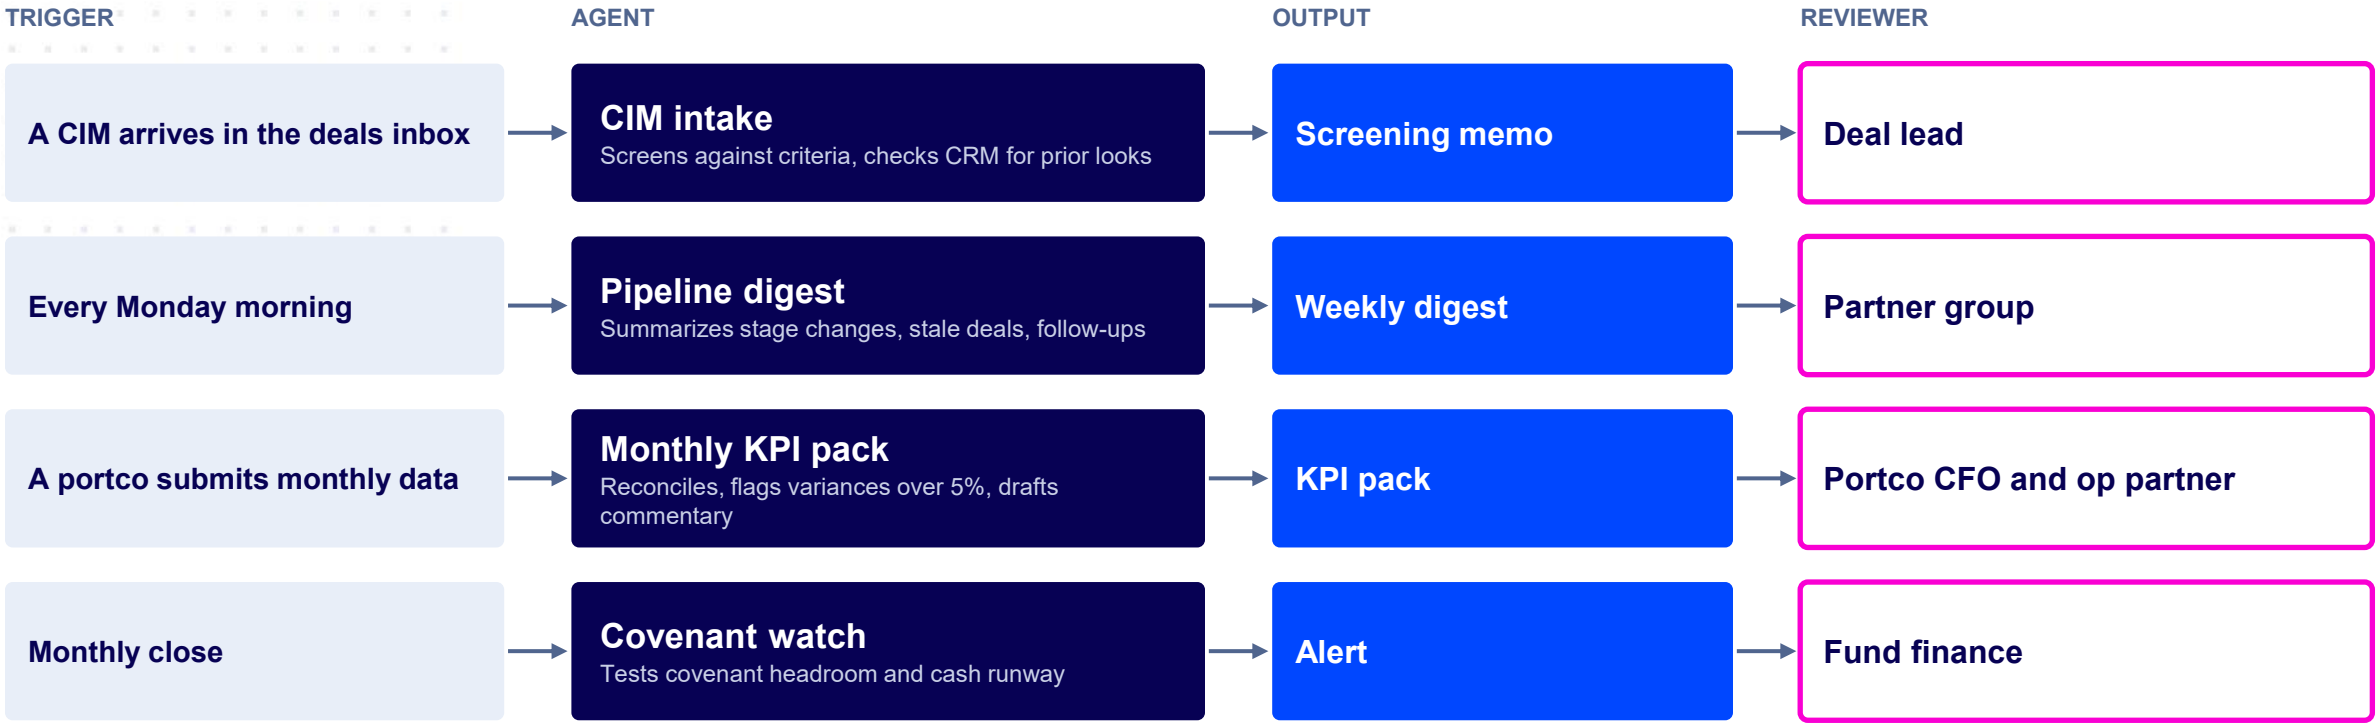
Con: Unstructured nuance + needs upkeep

Build Time: Months

1. **Start with RAG**
   over approved IC memos and diligence
2. **Add a context graph**
   companies, deals, and people
3. **Capture decision traces**
   from IC and pass notes

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

Recurring work moves to scheduled agents that report to a reviewer; agents being “always on” is starting to become more common (e.g, Dots)

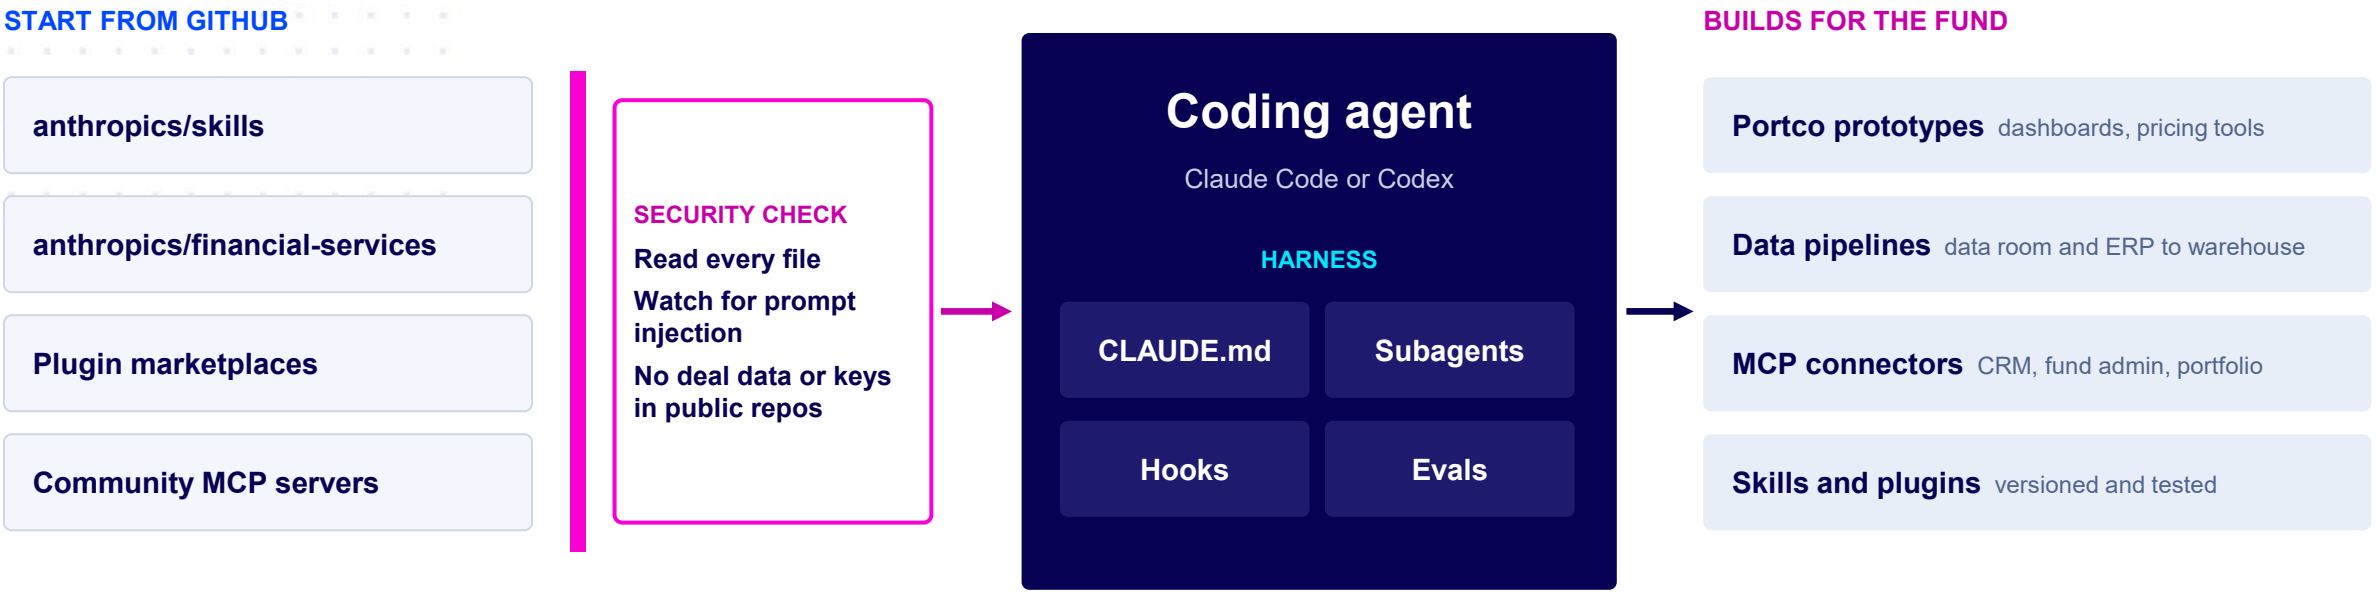


HOW THEY RUN

1. **Built** — by the pod, with test cases
2. **Deployed** — as managed or scheduled agents
3. **Connected** — through firm MCP servers
4. **Logged** — inputs, outputs, reviewer, cost

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

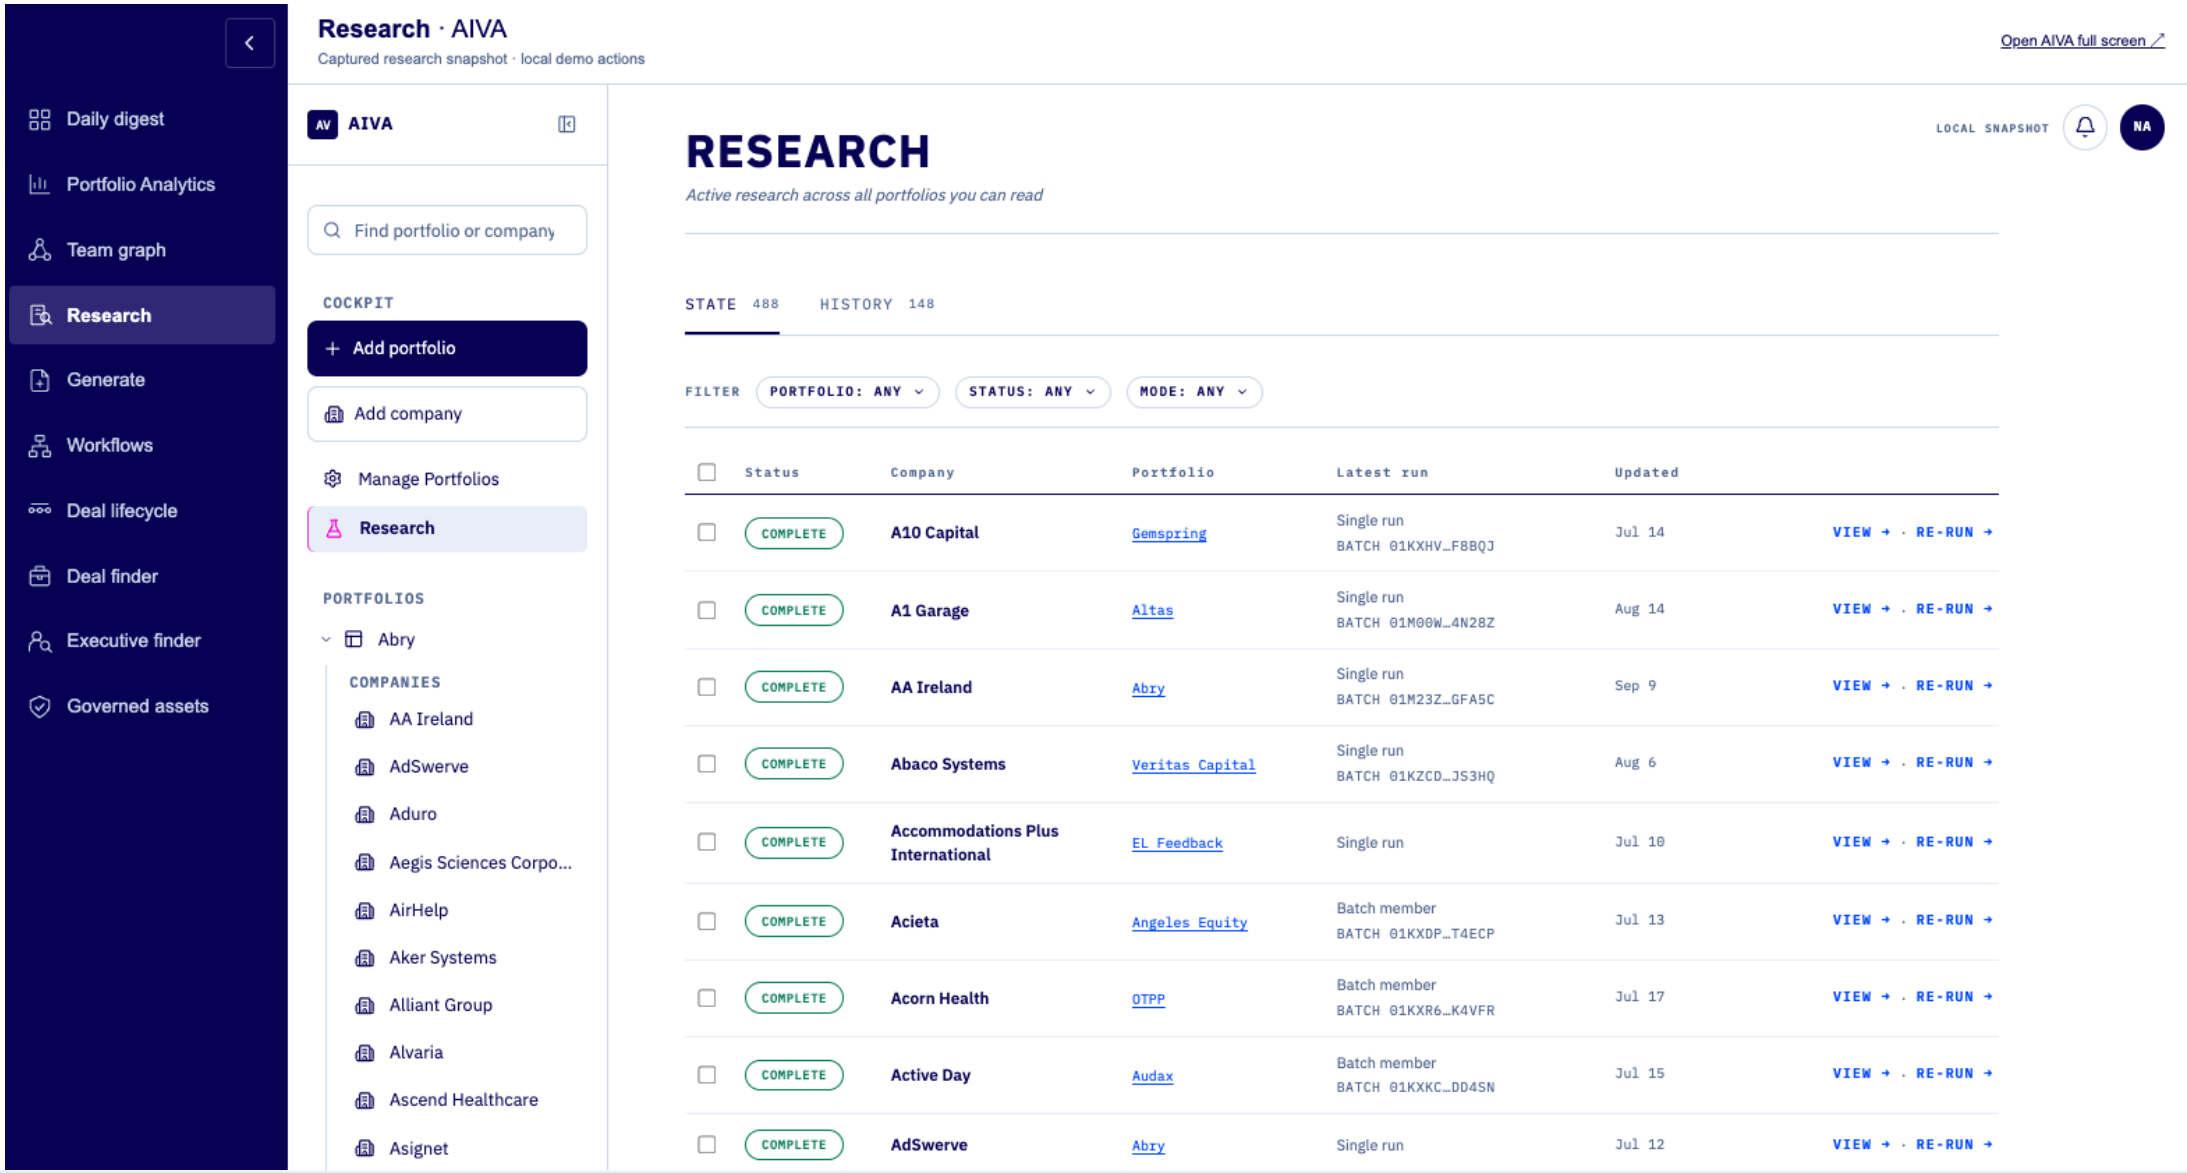
Coding agents build prototypes in days, starting from open skills on GitHub



Why it matters: a fund no longer needs an engineering team to try an idea. An FDI can stand up a working prototype on sample data in days, and the harness keeps it consistent and tested (good note here – emphasize E2E – end-to-end testing whenever building)

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

”Run" Demo



© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

Organized data unlocks use cases that prompts over attached files cannot reach

SOURCING

Relationship-led sourcing

Rank bankers by relevant deal flow sent to the fund and surface warm paths to targets

Needs: CRM, deal history, people graph

DILIGENCE

Pass-rationale analysis

Compare a new deal with past passes and losses, using the reasons recorded at IC

Needs: IC memos, decision log

ADVISORY

AI IC Committee Member

Build an AI IC committee member that takes in IC memos and provides perspective ahead of meetings

Needs: IC memos, historical data access

PORTFOLIO

Portfolio early warning

Flag KPI drift across portcos against plan and peer cohorts

Needs: KPI feeds, shared metric definitions, PortCo source integrations

FUNDRAISING

LP reporting from Source

Understand LP data and touchpoints from source systems and report to LPs

Needs: Fund admin data, approved answer library

EXIT

Exit evidence / VC pack

Trace realized value to initiatives and costs for buyer diligence based on source systems and trends

Needs: Value creation tracker, financials, source system integrations/MCPs

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Cost scales with usage, so the fund routes each task to the right model and sets budgets before agents run unattended

MODEL TIER	USE FOR	PE EXAMPLES
Frontier reasoning e.g., Claude Opus, GPT-6 Astra	Judgment, synthesis, long documents	IC memo drafts, thesis challenge, model review
Mid-tier e.g., Claude Sonnet, GPT-6/6.1	Most daily work and agent runs	Diligence extraction, KPI commentary, skill runs
Small and fast e.g., Claude Haiku, GPT-6/6.1 Luna	High-volume, simple tasks	Document tagging, CIM intake triage, email routing

WHERE THE MONEY GOES

Seats
Tens of dollars per user per month; about $100–$200 for heavy agent users

API usage
Metered per token for agents and scheduled jobs

People and platform
At Run, the pod and warehouse compute outweigh licenses

CONTROLS

Budgets
Set per team and per agent, with alerts

Routing
Default to the cheaper tier; escalate on failure

Caching and batch
Cut cost on recurring, repeatable runs

Unit cost
Track cost per accepted output each month

Seat figures reflect published list prices for business and premium plans as of September 2026; enterprise pricing is negotiated.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Next 90 days: move one step up from wherever your fund starts today

	STARTING AT CRAWL	STARTING AT WALK	STARTING AT RUN
Days 0–30	Approve enterprise tools and a one-page policy; connect PitchBook or Capital IQ	Pick two workflows with a baseline, such as the CIM screen and the KPI pack	Inventory data sources and owners; choose the first graph entities
Days 31–60	Assign the eight analyst tasks and log time savings and track improvements before and after	Hire or appoint a Forward Deployed Investor; set review gates by tier	Stand up firm MCP servers over CRM and deal data
Days 61–90	Turn the three best prompts into shared skills; name owners	Publish a firm skill catalog; scope the context layer on existing data	Launch two managed agents with test cases and cost budgets

Don’t forget to measure. Evaluating hours saved and benchmark improvements at Crawl, Walk, and Run will help you make informed decisions about where to invest further in AI across your firms.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

Key takeaways

Wherever you are on Crawl, Walk, or Run, these always matter…

YOUR NEXT STEP, BY STAGE

- **Crawl** Pick one recurring task and figure out what works best
- **Walk** Turn the best prompt into a shared skill
- **Run** Govern the data foundation and backend the agents run on

01

Keep the thinking human

Don’t outsource judgment, creativity, or problem-solving. It’s more important than ever.

02

Context beats clever prompts

AI only knows what it sees: give it the right files, definitions, directions, and firm data.

03

Verify before you trust

Cite every source, tie numbers to cells, keep a person on anything vital.

04

Make what works repeatable

Turn good prompts into shared skills, so everyone starts with a leg up.

05

Build, test, and iterate

It has never been faster to test ideas and solutions out. Have fun and figure out what works best for your firm.

06

Constantly rethink your AI framing

As everyone is aware, the rate of change is phenomenal and hard to keep up with. What was not possible 3 months ago may now be readily solved.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

Q&A

Questions and discussion

Thank you for joining!

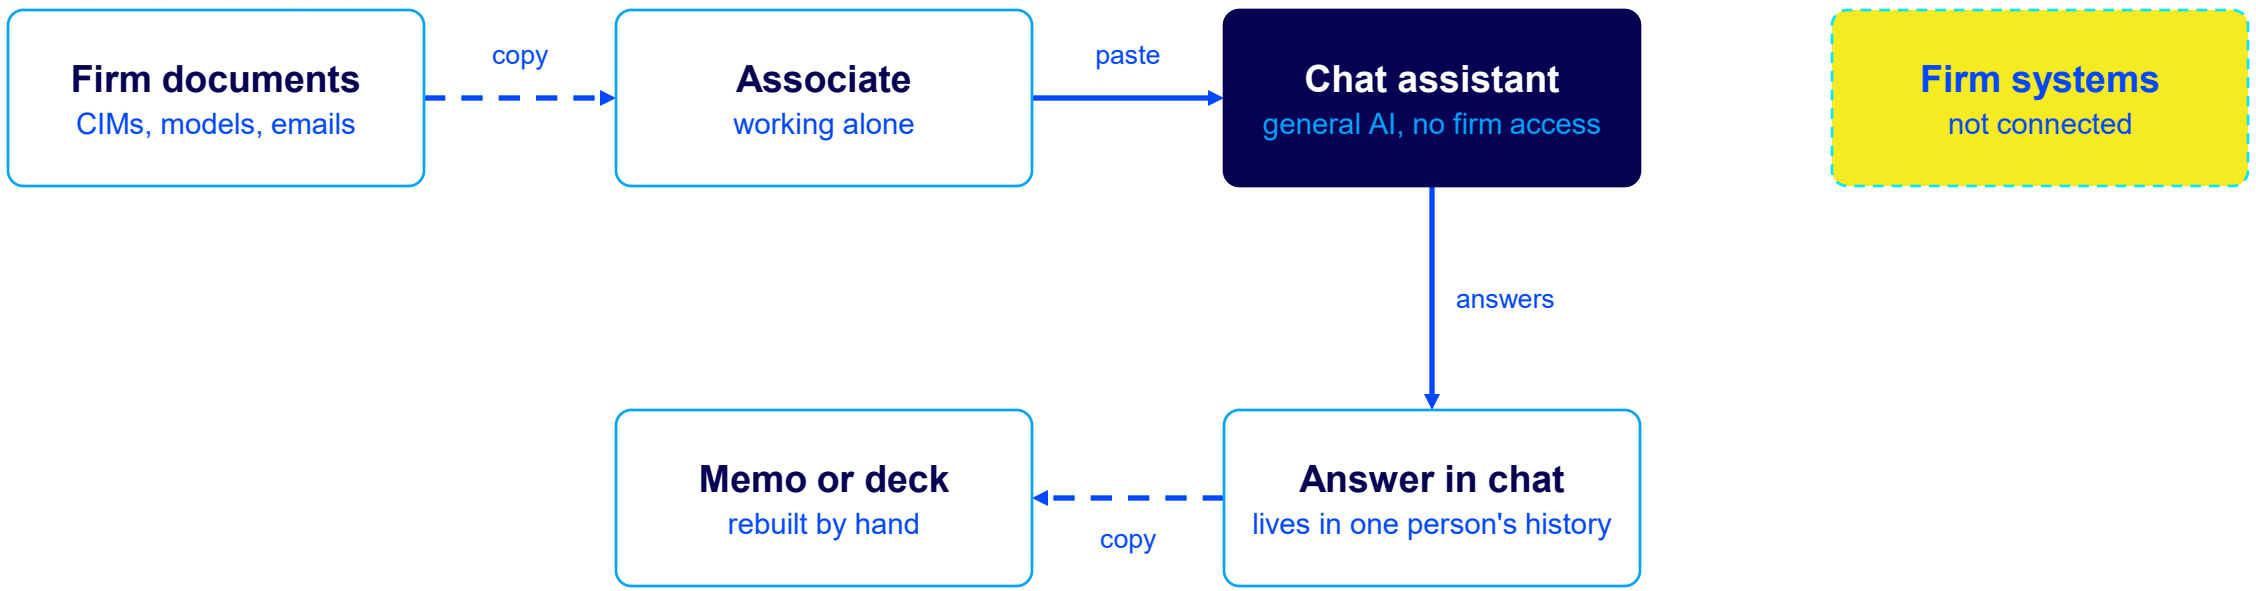
© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

APPENDIX

A	**Crawl, Walk, Run Diagrams** How the tech flows at each stage	39–41
B	**AI Tooling in Finance** The landscape across the deal lifecycle	42
C	**Stall Points + Resolution** Why AI stalls, benchmarking, mobile access	43–46
D	**Regulation in PE AI** SEC and other rules for PE firms	47
E	**Context Engines** Horizontal, PE-specific, or built	48
F	**Caching and Context Windows** How prompt caching and context windows work	49–50

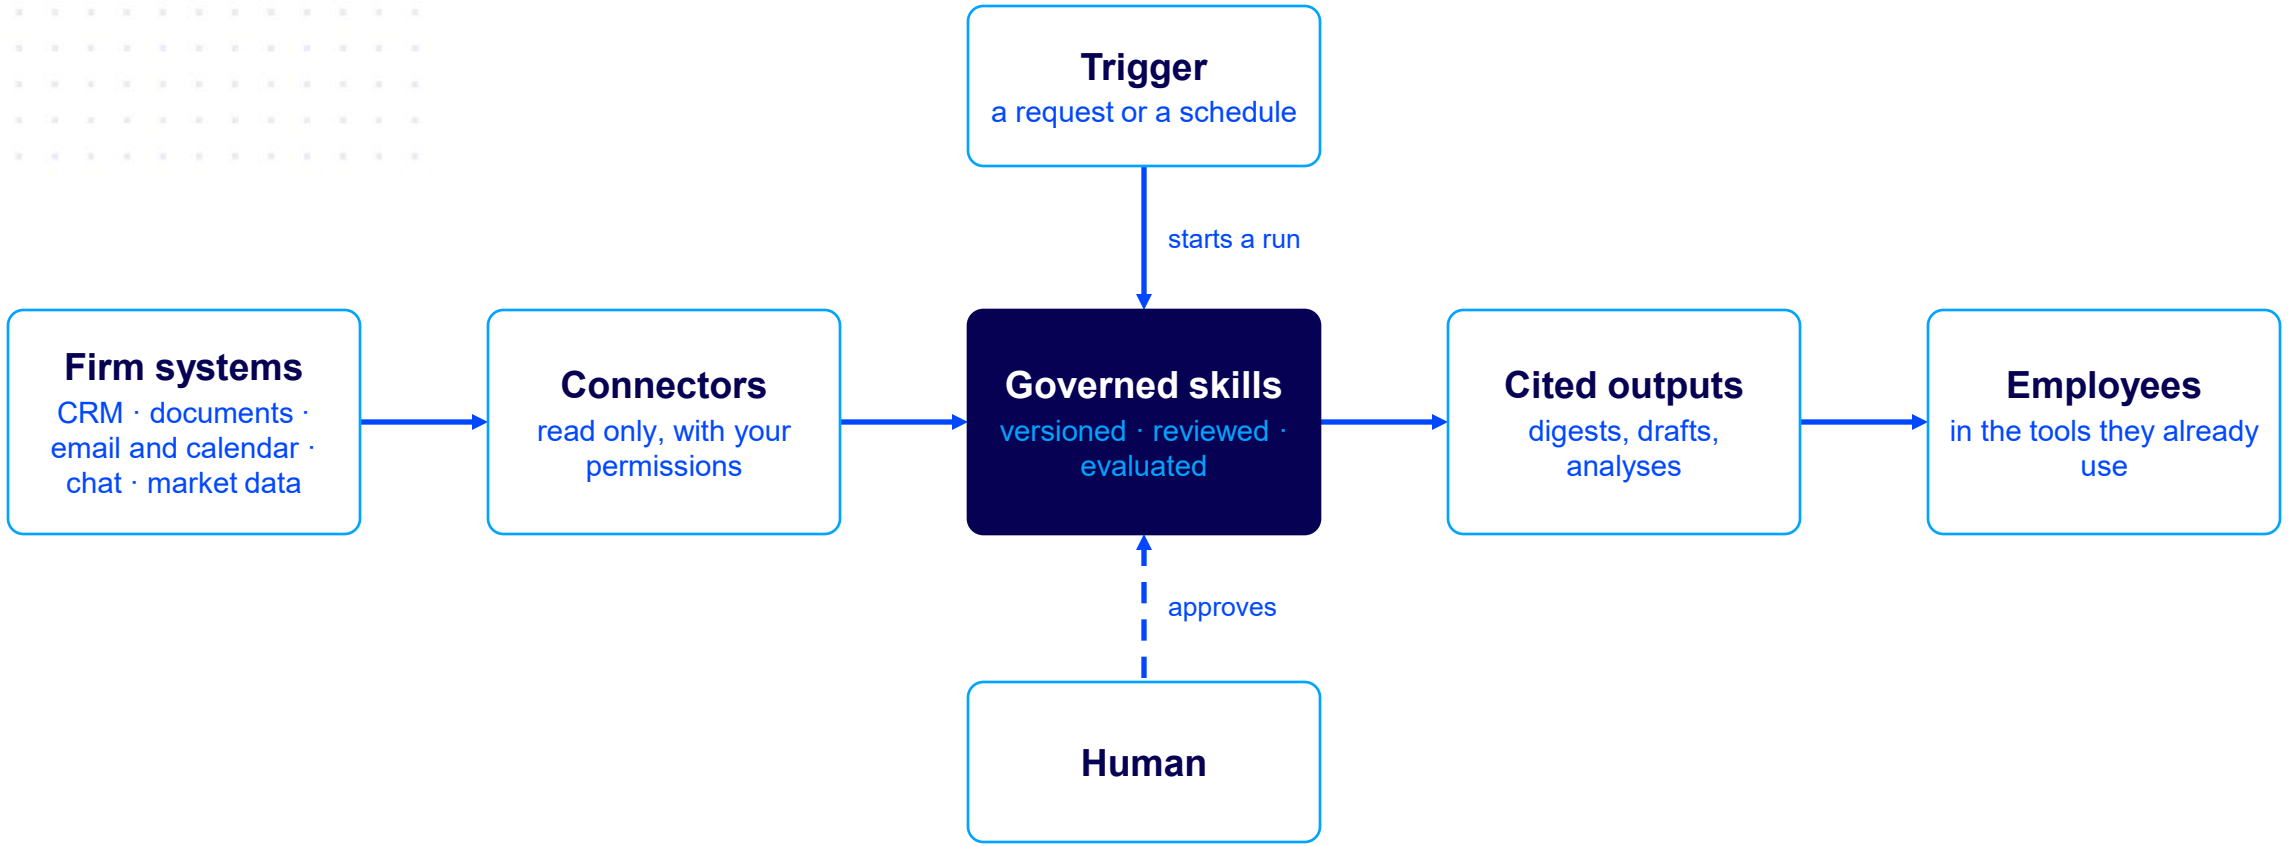
© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Copy in, copy out



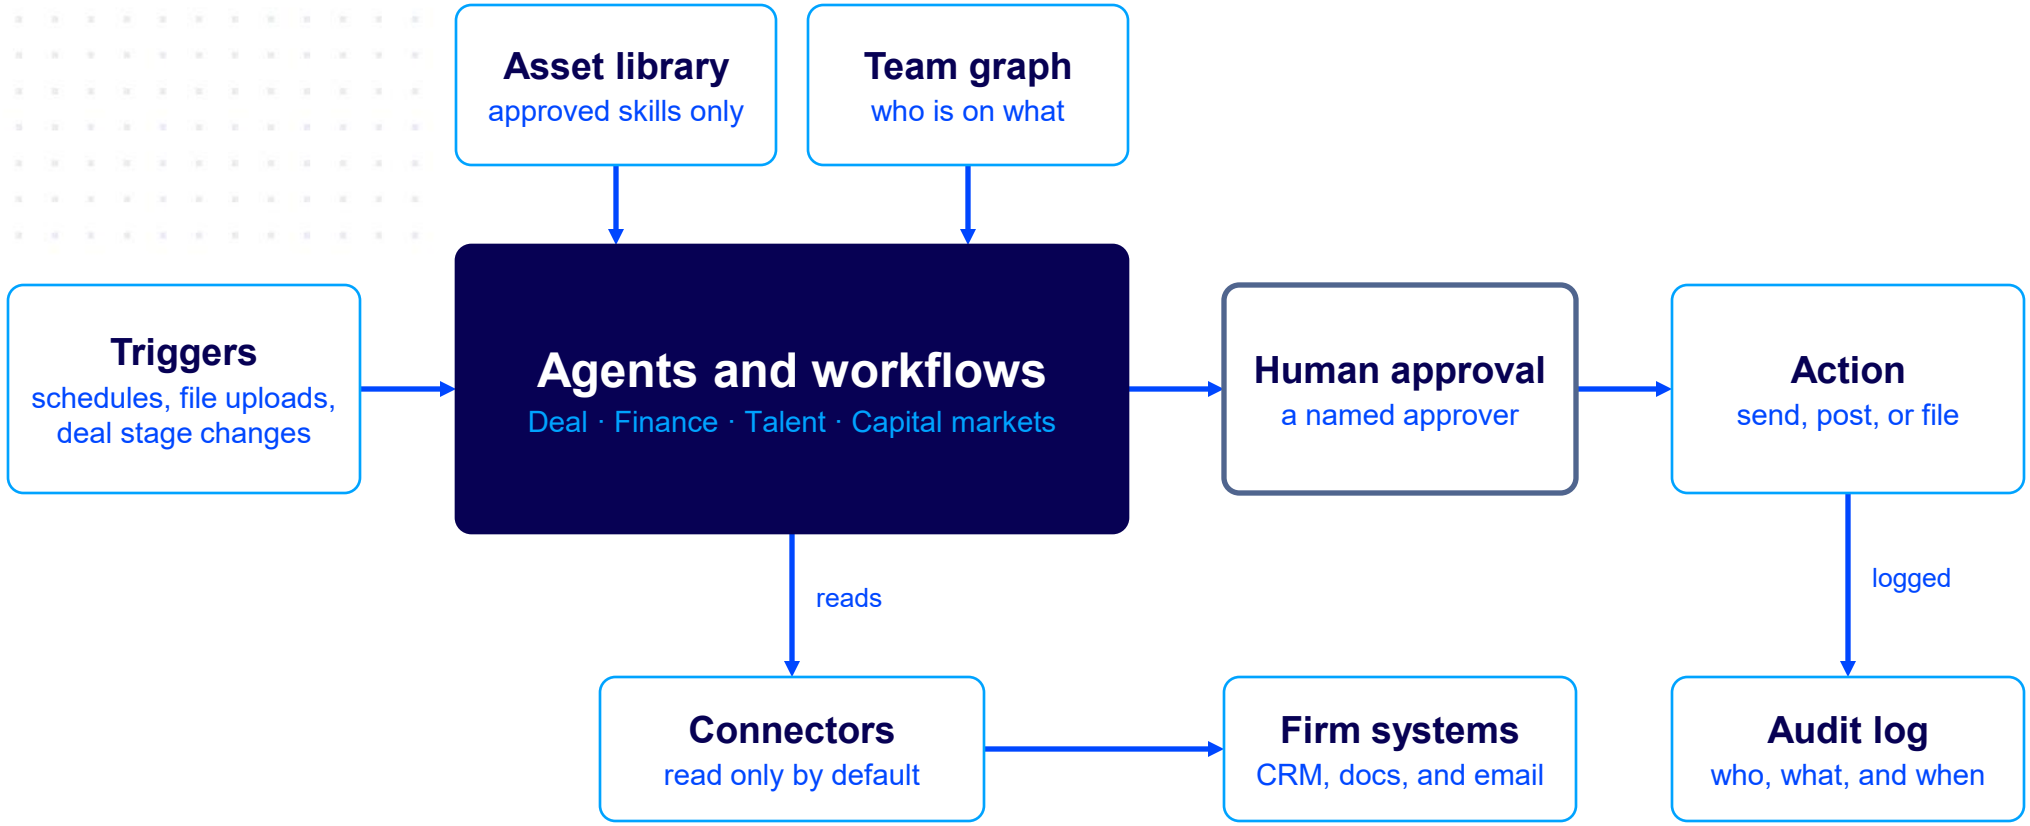
© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Reviewed skills, connected to firm systems



© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Agents do the prep, people make the call



© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Latest financial AI tooling can cover the entire deal lifecycle

Sourcing	Research and diligence	Data and data rooms	Portfolio monitoring	Fund ops and IR
Find targets and warm paths	Read, compare, and draft	Licensed data and VDRs	KPIs, valuations, covenants	Accounting, LPs, DDQs
Grata	AlphaSense	S&P Capital IQ	Chronograph	Allvue
SourceScrub	Hebbia	FactSet	iLEVEL	Juniper Square
PitchBook	Rogo	Morningstar	Cobalt	Carta
DealCloud	Blueflame AI	Daloopa	73 Strings	Canoe
Affinity	Keye	Datasite	Lumonic	
	Transacted	Intralinks	Standard Metrics	

GENERAL-PURPOSE AI THAT WORKS ACROSS EVERY STAGE

Claude
Excel add-in; connectors to Capital IQ, PitchBook, Morningstar, Daloopa; finance skills

ChatGPT
Enterprise workspace, agents, connectors to firm files

Microsoft 365 Copilot
AI inside Excel, PowerPoint, Word, Outlook, Teams

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Eight reasons AI stalls trace back to three root causes, each with a known fix

No firm context

The model guesses because it cannot see the firm

Fixed at Walk

It makes things up

→ Ground it in firm data, cite every claim, keep math in formulas

Every answer is different

→ Reviewed skills, evals before release, fixed formats

It doesn’t know our firm

→ Connect core systems; put house definitions in skills

No governance

Nobody owns, reviews, or records AI use

Fixed at Walk

We can’t put our data in it

→ Enterprise terms, inherited permissions, info barriers

Nobody knows who uses it

→ Asset library, AI Governance review, usage reporting

We won’t let it act alone

→ Agents draft, people send; full audit trail

No operating model

Good results never become how the firm works

Fixed at Run

Pilots never spread

→ Save as a workflow, a review path, team champions

We can’t show the value

→ Pick recurring work, baseline first, track reuse

Context and governance arrive at Walk. The operating model matures at Run.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners’ prior consent is prohibited.

Fix the causes in order

Crawl
Individuals experiment

Enterprise accounts

A usage policy

Walk
Governed and connected

Read-only connectors

Citations on every output

Reviewed skills and evals

Permissions and barriers

Run
Agents across the firm

Approval for anything external

A full audit log

Save-as-workflow and review

Value tracking by team

THE FIRST 90 DAYS

1

DAYS 1–30

Pick and baseline

Three recurring workflows; record today's hours

2

DAYS 31–60

Connect and govern

Read-only connectors and an AI Governance review

3

DAYS 61–90

Ship and measure

First reviewed skills live; report against baseline

Build the controls first. Trust follows.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Benchmark every skill, workflow, and context engine on your own deals

THE BENCHMARK LOOP

1. **Build a golden set**
   Past deals with known answers, plus hard cases
2. **Set a baseline**
   Today's manual process or current version
3. **Run and score**
   Automated checks plus a blind expert rubric
4. **Set the release bar**
   Minimum scores to publish
5. **Rerun on changes**
   New model, new data, or new version

↻ repeat on every change

Skills
Does it do the task right?

- Figures tie to source
- Citation coverage
- Format followed
- Reviewer edits needed

Workflows and agents
Does it finish the job?

- Completion without rework
- End-to-end cycle time
- Escalation and failure rate
- Cost per accepted output

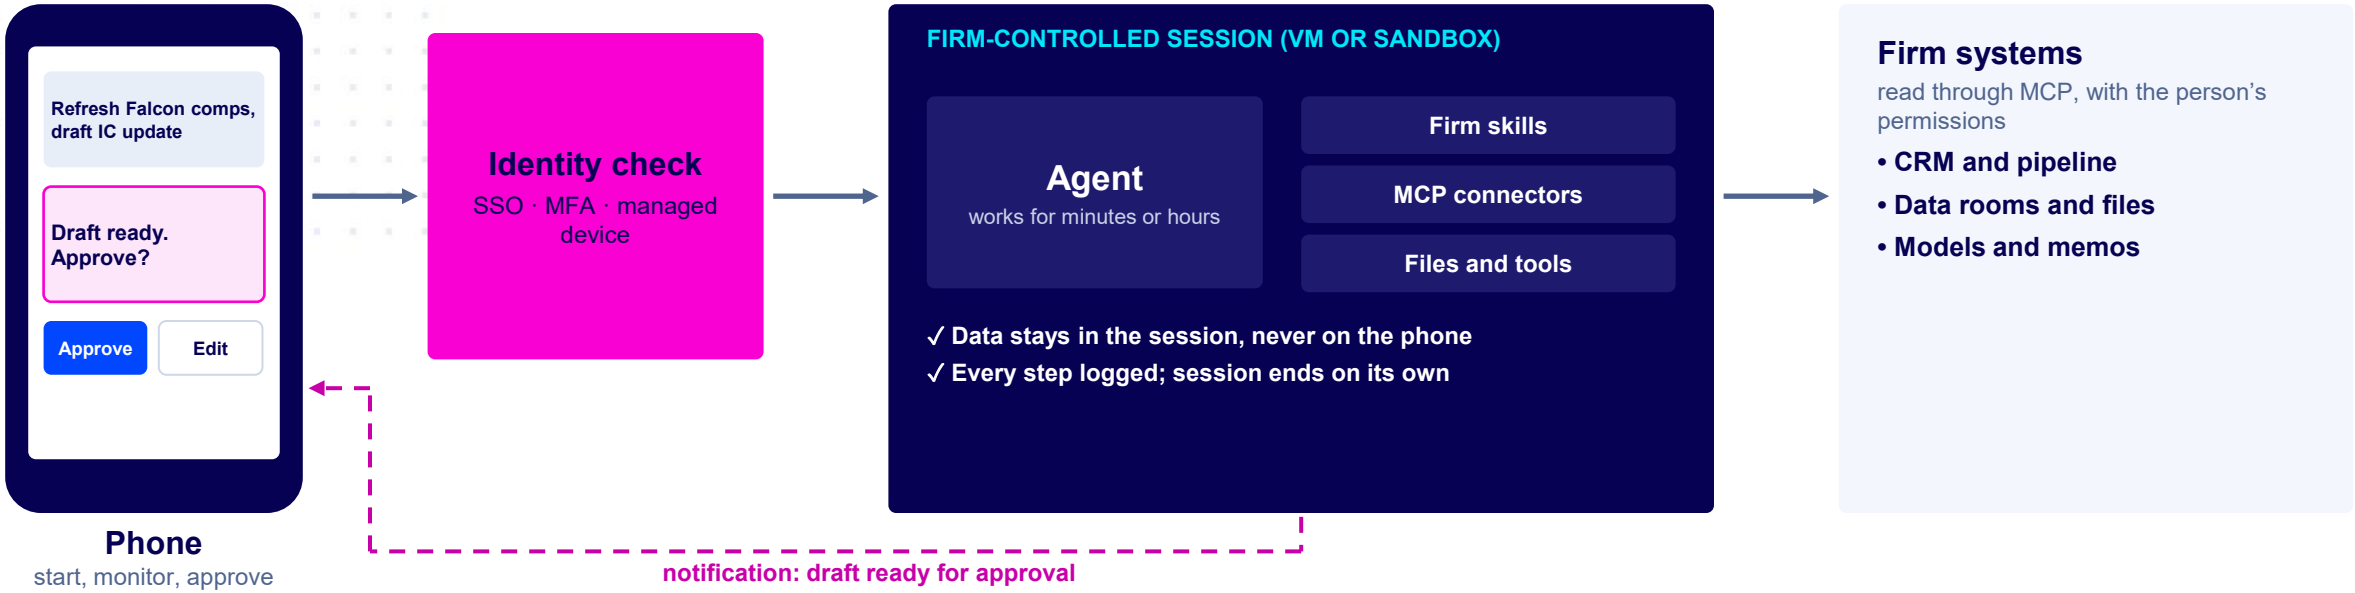
Context engines
Does it find the right thing?

- Right source retrieved
- Answers grounded in sources
- Permissions respected
- Freshness of data

RULES OF THUMB: Use your data (not vendor leaderboards), include missing-data and conflicting-source cases, and score blind against the human baseline

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Agents run in firm-controlled sessions; people approve from their phone



Hosted agent sessions
The AI vendor runs the session; you reach it from its mobile app
Fastest to start; check data terms

Firm cloud desktops
Virtual desktops where desktop agents run, reached from any device
Uses controls IT already has

Firm-hosted sandboxes
Containers in your own cloud tenant, triggered from chat or mobile
Most control; needs an engineer

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

There is no AI-specific SEC rule today; existing rules already apply to AI use

SEC: RULES THAT ALREADY COVER AI

Fiduciary duty (Advisers Act)
IN FORCE

AI-assisted advice and disclosures must meet the duty of care and loyalty.

Marketing Rule and anti-fraud
IN FORCE

"AI washing": Delphia and Global Predictions settled for $400K combined (Mar 2024) over false AI claims.

FY2026 exam priorities
ACTIVE

Examiners review the accuracy of AI claims, whether policies monitor and supervise AI use, and AI-related cyber threats.

Regulation S-P amendments
IN FORCE

Incident response, breach notice, and service-provider oversight. Compliance: Dec 3, 2025 (≥$1.5B AUM); Jun 3, 2026 (smaller).

NOT IN FORCE

Predictive data analytics rule

Proposed 2023 to govern AI-driven conflicts. Withdrawn Jun 12, 2025; the SEC says any future action needs a new proposal.

Federal preemption of state laws

EO 14365 (Dec 2025) and the White House framework (Mar 2026) push Congress to preempt; neither is law yet.

BEYOND THE SEC (PORTCOS AND EU)

EU AI Act
PARTLY IN FORCE

GPAI duties since Aug 2025; transparency (Art. 50) since Aug 2, 2026. High-risk deadline moved to Dec 2, 2027.

Colorado AI Act (SB 189)
2027

Narrowed to disclosure and human review for automated decisions; effective Jan 1, 2027.

Other state laws
IN FORCE

Remain in force until Congress preempts them or courts strike them down.

Not legal advice. Sources: SEC FY2026 Examination Priorities; SEC Release 33-11377 (Jun 12, 2025); SEC v. Delphia / Global Predictions (Mar 18, 2024); Reg S-P amendments (May 2024); EO 14365; White House AI Framework (Mar 20, 2026); EU Reg. 2026/1744; Colorado SB 189.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Context engines make your AI firm-aware: horizontal, PE-specific, or built

AI YOUR TEAM USES

Claude | ChatGPT | Microsoft Copilot | Gemini | Your own agents

MCP · one open standard lets any of these AI tools read the context layer, with each person's permissions

Horizontal enterprise
Knowledge across every SaaS tool

- Glean
- Microsoft IQ
- Atlassian Rovo
- Gemini Enterprise

Private-capital specific
Deals, documents, and decisions

- Metal
- Chronograph (MCP)
- Lumonic (MCP)

Build your own
Your platform, your schema

- Snowflake / Databricks / Fabric
- Vector DB + graph
- Firm MCP servers

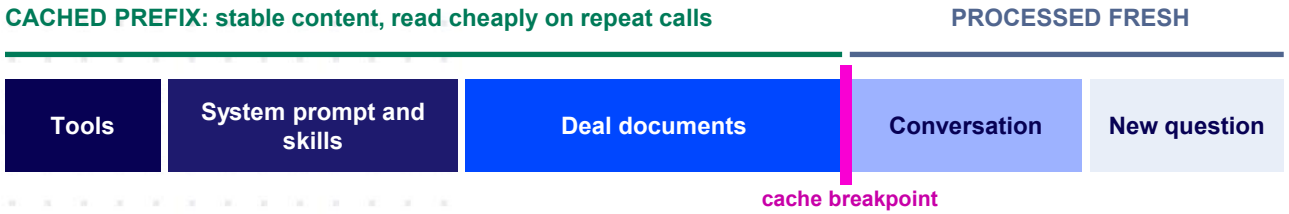
YOUR SOURCES

CRM | SharePoint / Drive | Email and Teams | Data rooms | Portfolio data | Licensed market data

Favor context engines/platforms that expose your context through MCP and APIs.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

Caching reuses the unchanged front of a prompt, so repeats are cheaper



CACHE HITS AND MISSES



DESIGN FOR HITS

- Stable content first, changing content last
- No timestamps or IDs at the top
- Keep tool and skill definitions fixed

OTHER CACHES

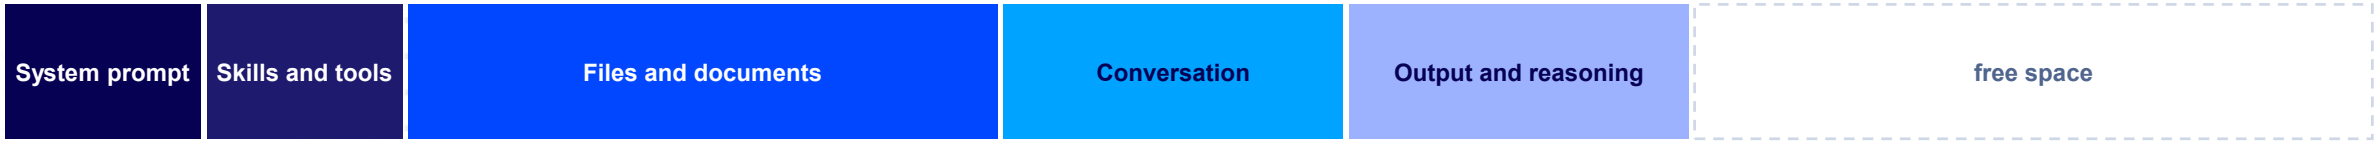
KV cache: inside the model, within one response. Semantic cache: your app reuses answers to similar questions; watch for stale answers.

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.

A context window is the model's working memory: all it can see at once

WHAT FILLS A 1M-TOKEN WINDOW (ILLUSTRATIVE)

1M tokens ≈ 555K words



Bigger is not always better

Recall gets less reliable as the window fills. Long, mixed sessions are where errors creep in.

A window is not memory

It resets with every new session. Anything the AI should know must be loaded again, or live in a skill or connector.

Input and output are separate limits

Current Claude models read up to 1M tokens but write up to 128K in one reply.

KEEP THE WINDOW CLEAN

One task per session
start fresh with a short summary

Load only what is needed
skills and connectors fetch on demand

Compact long sessions
summarize older turns

Split big jobs
subagents each get their own window

© 2026 West Monroe Partners | Reproduction and/or distribution without West Monroe Partners' prior consent is prohibited.