Juniper Square

Making Your Firm AI-Legible

A GP's Guide for Putting AI to Work

Contents

Introduction: The operational alpha of AI 3

Part 1: A six-step playbook for AI transformation 4

1 Pick your first workflow 6

2 Define what good looks like 9

3 Figure out what AI needs to know 13

4 Turn knowledge into reusable context with substrates 16

5 Give AI access to data and tools 23

6 Decide how to implement the workflow 31

Part 2: Top workflows for AI transformation 40

Part 3: A primer for executives managing AI initiatives 46

Conclusion: Your data is the head start 55

The operational alpha of AI

Most GPs are already using AI, but the gap between what they expect and what they get is often wide. Many also assume every other firm is already deep into AI adoption. In practice, most are still near the starting line, usually for structural reasons.

We see the cost every day: teams use AI for simple chat and drafting use cases, but when it comes to getting real work done, they fall back on the manual processes that they've spent years developing. Teams that are asked to "adopt AI" but not given a clear roadmap struggle, and tension can build inside firms as leaders don't see the progress they expect to happen.

In reality, every firm already has the most important ingredient for AI success, which is its operating model: the knowledge, processes, data, and judgment behind how it raises capital and runs its funds. Much of it is scattered across systems or lives in people's heads, out of AI's reach. The firms that will gain the most will make that knowledge accessible, then redesign their workflows around what becomes possible.

For GPs, the opportunity with AI goes well beyond making existing work faster. It's a chance to rethink how work gets done across the firm, and, as a result, generate more operational alpha.

> **Juniper Square's CEO, Alex Robinson, puts it this way:**
>
> "AI is going to force you to reshuffle how all the work is done inside your organization. It requires being bold, and operating with imperfect information, but the core of what you're trying to do is to make your company AI-legible."

This guide is a practical playbook for making your firm AI-legible.

PART 1
Six steps for building a useful AI workflow.

PART 2
Proven AI use cases for private markets, and how our customers are putting them to work.

PART 3
Advice for executives on how to lead teams through AI initiatives.

PART 1

A six-step playbook for AI transformation

We've found that the best way to move beyond individual AI usage is to take one meaningful, firm-wide use case and build it out. This takes AI transformation from an opaque buzzword into a practical mission teams can reasonably drive to success, while still managing their day-to-day responsibilities. This playbook will take you on the full journey, from identifying where to start to being ready to implement. All of this happens in six steps:

1 **Pick your first workflow** p. 6

2 **Define what good looks like** p. 9

3 **Figure out what AI needs to know** p. 13

4 **Turn knowledge into reusable context with substrates** p. 16

5 **Give AI access to data and tools** p. 23

6 **Decide how to implement the workflow** p. 31

Before diving in, know that the tools will change quickly. Models will improve and new ways of building agents will emerge, so some guidance today may look different in a year. However, the core concepts will hold: AI needs a defined job and the right data and context, along with instructions for applying your firm's judgment and guardrails on what it can do.

TRY IT

Determine your firm's AI readiness

Before you start the playbook, assess your firm's AI readiness so the prep work is done and your AI transformation has a better shot at success. At Juniper Square, we've identified eight enablers your firm needs to be AI-ready, as part of our AI assessment:

- **Use cases**
- **Operating model**
- **Economics**
- **IP & knowledge**
- **Systems**
- **Data & contracts**
- **Talent**
- **Governance & security**

Want to know how ready your firm is for AI? Take our assessment to get a detailed report on where your firm stands today and what you should do next to improve your AI readiness.

1

STEP 1 OF 6

Pick your first workflow

For many firms, the best AI use case emerges from an operational hurdle that can no longer be ignored: manually looking across fragmented data sources to build a fund rollup, doing yet another complex reconciliation in a bespoke spreadsheet, or repeatedly answering investor questions that require pulling together the same context.

Start with a problem that's valuable enough to matter and painful enough that changing the way it gets done would make a noticeable difference. If you're not sure where to start, the examples in Part 2 may provide inspiration See Part 2, p. 40 →, or you can use these questions to help narrow your options. The more of these you can answer yes to, the better the starting point.

- Does this process or workflow happen often enough to matter, rather than just once a quarter?
- Is there enough data or context already available for AI to work with?
- Is this more about following a pattern than exercising judgment or creativity?
- Is this a process that's already working well, where the team agrees on what "good" looks like?
- Does the team have the capacity to build and test something here, or is everyone just keeping the lights on?

TIP

Don’t just automate your current process

When electric motors first reached American factories, manufacturers often used them as direct replacements for the steam engines that powered their existing machinery. But the factories themselves were still organized around a central source of power, distributed through shafts, belts, and pulleys. The bigger gains came later, when manufacturers realized that smaller electric motors could power individual machines directly. That allowed them to redesign factories around the flow of production rather than the constraints of a central engine.

AI creates a similar opportunity. Once you’ve identified a workflow worth improving, don’t assume the goal is to reproduce every existing step with AI. Ask what the workflow would look like if you redesigned it based on what AI can do today. Some steps may disappear, others may combine, and work that previously had to happen sequentially may no longer need to.

WORKSHEET

Score your candidate workflows

Use these questions to help narrow your options. The more of these you can answer yes to, the better the starting point.

Question	Workflow A	Workflow B	Workflow C
Workflow name			
Does this process or workflow happen often enough to matter, rather than just once a quarter?	☐	☐	☐
Is there enough data or context already available for AI to work with?	☐	☐	☐
Is this more about following a pattern than exercising judgment or creativity?	☐	☐	☐
Is this a process that's already working well, where the team agrees on what "good" looks like?	☐	☐	☐
Does the team have the capacity to build and test something here, or is everyone just keeping the lights on?	☐	☐	☐
Total yes answers			

Ask what the workflow would look like if you redesigned it based on what AI can do today.

Steps that may disappear	Steps that may combine	Work that no longer needs to happen sequentially

2

STEP 2 OF 6

Define what good looks like

Once you have a use case in mind, get specific about the outcome you want.

Imagine you want to design an investor tearsheet with AI. Start by describing your ideal tearsheet. Do you want to show high-level financial information (e.g., total committed capital) or more granular details (e.g., individual capital calls or fund-level performance)? How much information do you want to add about the investor themselves (e.g., contact information, overview of investment mandate, notes from past interactions)? Do you want it to look branded or more generic? You can also pull tearsheets you've recently built to provide concrete examples of the output you're looking for.

TRY IT

Before you build, write down what you're asking AI to accomplish using a workflow brief.

The brief does two jobs: 1) it aligns your internal stakeholders on what good looks like and sets the bar for consistency and quality, especially if several people use the workflow; 2) it gives you the success criteria and guardrails you'll translate into instructions for AI, and shows the data and context the workflow needs.

Use your answers to scope the workflow and get alignment with your broader team. Keep the first version simple, and don't worry yet about how AI will produce the result.

Example and worksheet on the next pages →

EXAMPLE

A completed workflow brief

PROBLEM

1 What are you trying to improve or change? Why does it matter?

Before every LP meeting, an IR associate spends a few hours pulling the investor's commitments, capital calls, and recent emails from three systems into a one-page tearsheet. Quality varies from person to person, and the prep crowds out time for the conversation itself.

2 Who will use the output, and what will they do with it?

The IR associate and the partner leading the meeting. They use it to prepare talking points and spot open follow-ups before they walk in.

3 What should the workflow not do?

It should not change any data in the CRM or fund accounting platform, and it should never send the tearsheet to the investor.

OUTPUT

4 What should the workflow produce?

A one-page branded PDF with four sections: investor profile, commitments and called capital by fund, recent correspondence, and open follow-ups. Financial figures come from the fund accounting platform; contact details and notes come from the CRM.

5 What information would the AI need to do this task well?

Investor positions and transaction history, CRM contact records and meeting notes, recent email threads, our definitions for committed and called capital, and three recent tearsheets the team considers strong.

6 When should a human be in the loop?

AI drafts the tearsheet; the IR associate checks every figure before it's shared. If two systems disagree on a number, AI flags it instead of choosing one.

7 When does this workflow need to run?

Automatically two days before a scheduled LP meeting, or on request.

SUCCESS CRITERIA

8 How will you know the result is good?

Every financial figure matches the system of record, and the IR team can use the draft without rebuilding it for at least nine out of 10 meetings.

9 What does strong work look like?

The tearsheets the team prepared for last year's annual meeting.

WORKSHEET

Workflow brief

PROBLEM

1 What are you trying to improve or change? Why does it matter?

Describe how the work happens today, where the friction is, and why improving it matters. Focus on the business problem rather than the technology.

2 Who will use the output, and what will they do with it?

Identify the primary user and the decision or next action the workflow should support.

3 What should the workflow not do?

Spell out important exclusions and failure conditions. For example, the AI should not change data in the system of record or send any investor communications automatically.

OUTPUT

4 What should the workflow produce?

Describe what the final result should contain and how it should be structured. Be concrete. For instance, should the output be a report, an alert, a dashboard, a Slack message, or something else? What sources should the AI use in generating its output, and how should it compute its ranking, calculations, or summary?

5 What information would the AI need to do this task well?

List the types of information that the AI should use to inform the output, and, if you know, which system or source that information might come from. Examples of information include fund and investment data, CRM activity, meeting transcripts, internal policies or definitions, best practices, examples of completed work, etc.

WORKSHEET

6 When should a human be in the loop?

Define where human judgment remains necessary. Consider: Can the workflow make recommendations, or should it only surface information? Can it take actions automatically, or are there steps that require approval? When should AI escalate for human review?

7 When does this workflow need to run?

Perhaps the workflow is triggered by an event, or it runs on a recurring schedule. Spell that out here.

SUCCESS CRITERIA

8 How will you know the result is good?

Consider accuracy, completeness, consistency, timeliness, usefulness, format, and anything else that matters for this task. Examples: All underlying financial figures match the system of record. At least 90% of known investor relationships are correctly identified.

9 What does strong work look like?

Include a few examples of high-quality past outputs, analyses, reports, or decisions if they exist. These give AI—and the people building the workflow—a concrete reference for the standard you are trying to meet.

Your tolerance for error is also a clue about which use cases to start with. Some of the most consistently successful ones we see firms using so far are synthesis use cases (e.g., pulling together scattered information and making sense of it) and notification and alerting use cases (e.g., surfacing things worth a second look), precisely because a partial result still adds value. A workflow that has to be right every time is a much higher bar to clear.

STEP 3 OF 6

3 Figure out what AI needs to know

Once you know what you want AI to produce, work backward: What would it need to know to do this well?

Some of that information will already live in systems of record, such as your CRM, general ledger, fund accounting platform, or investor management system. But some of the most important context may live elsewhere—or only in someone's head.

For an investor tearsheet, for example, AI may need to know where to find an investor's positions and transaction history, but also which information belongs on the tearsheet, how your firm defines particular metrics, which source to trust when two systems disagree, and how the output should differ for different types of investors.

DEFINITION

This is what we mean by making your firm **AI-legible**: capturing the business-critical information and judgment AI needs in an accessible format it can use.

Start with the knowledge your use case requires. Which systems hold the relevant data? Which meetings, emails, documents, or conversations contain useful context? What does an experienced employee know about this workflow that isn't written down anywhere?

Write down the fuzzier, more bespoke ways you operate, too: How does your company evaluate a deal? What separates a good investor from a bad one? Which deals do you chase, and which do you pass on? How do you source new deals? Capturing these processes helps an AI workflow reach the same conclusions that your team would have come to.

We've also found that by giving your AI tool six to eight good examples of the output you want along with two to four bad examples, you can ask it to give you a generalized best practice.

> "The models are [now] good enough to automate the vast majority of work that's happening inside a company," Alex Robinson, Juniper Square's CEO, explains. Given this reality, the foundational work firms must own is building the layer of durable knowledge that allows AI to work with the particularities of their business.

TRY IT

Get tacit knowledge out of someone's head

If an experienced employee knows how to do a task but has never documented their process, use AI to help extract it. You can ask ChatGPT or Claude to interview the person who performs the work. Give it the workflow brief you created in the previous step, then have it ask questions one at a time about how the person approaches the task—including the decisions they make, exceptions they encounter, and judgment calls that might not appear in a traditional SOP. ← See Step 2, p. 11

You can also give AI a recording or transcript of someone performing the task and ask it to identify the process and decision-making behind what they did.

Worksheet on the next page →

WORKSHEET

Interview someone who performs the work today

Task _____ **Person interviewed** _____

Have the person you are interviewing walk you through their workflow. At each key step, ask the questions below and make note of the answers.

How do you decide what to do next?

What information do you look at before making that decision?

What are the common exceptions to the usual process?

What would make you distrust the information you're looking at?

What mistakes would a new employee be likely to make?

When would you stop and ask someone else for help?

STEP 4 OF 6

4 Turn knowledge into reusable context with substrates

Once you've identified what AI needs to know, capture that knowledge in a form it can use repeatedly.

Some of that knowledge is context: durable information about your firm that may be useful across many tasks, such as what certain data fields mean, how entities map across systems, or how your firm talks about its investors. Other knowledge takes the form of instructions: how your firm performs a particular task, including the steps to follow, judgment to apply, and what a good result looks like.

Make both kinds of knowledge reusable so you aren't explaining them from scratch every time you run the workflow.

DEFINITION

At Juniper Square, we call shared knowledge that can form the foundation of multiple AI workflows "**substrates.**" A substrate holds knowledge you would otherwise have to give to the AI model with every request, such as what your data fields mean or how entities map across systems. Because many workflows draw on it, building a substrate once improves every workflow built on it. Instructions for performing a particular task become skills, which we cover below. See p. 20 →

Capture durable knowledge

Substrates give AI the firm-specific knowledge it needs to interpret your data and instructions correctly. Treat each substrate like a README for a particular part of your business. Write it as if you were explaining that domain to a smart new employee who lacks the institutional context your team takes for granted. A good substrate should cover:

- **Purpose and scope:** What knowledge or questions does this substrate cover, and what's out of scope?
- **Source and ownership:** Where does the knowledge come from, who owns it, and when was it last updated?
- **Core concepts and mappings:** Explain the important definitions and relationships AI needs to understand. For instance, distinguish committed capital from invested capital, or explain that "CalSTRS" in the CRM refers to the same entity as "California State Teachers Retirement System" in the fund accounting platform.
- **Maintenance:** When should the substrate be updated, and what events should trigger a review?
- **Handling disagreement:** If two sources conflict, which one should AI trust?

On the next page you'll find an abridged sample substrate we built at Juniper Square to help one of our clients understand their internal data sources.

EXAMPLE

Knowledge sources

This file combines knowledge from several systems and documents. Not every source is connected live, so use the information below to understand where each type of knowledge comes from and how current it is.

Knowledge	Canonical source	Internal owner	Status
Investment criteria & deal process	Investment Process (Markdown file)	Investment Team	Reviewed semi-annually
Prospective deals	Deal CRM	Investment Team	Live, available via MCP
Portfolio KPIs and operating data	Portfolio monitoring system	Investment Team	In implementation
Accounting transactions & portfolio valuations	Juniper Square	Finance	Live, available via MCP
Investor activity & biographical data	Juniper Square	Investor Relations	Live, available via MCP
Recorded calls (deals & investors)	Meeting notetaker	IT	Live, available via MCP

When two sources disagree, defer to the designated canonical source for that type of information. Each system of record should have an associated Markdown file (the substrate document) that explains what its data means and how it relates to other systems, so AI outputs built on it stay reliable.

Maintaining your substrates over time

A substrate is only useful while it stays accurate. Assign each substrate an owner, and write the five rules below into the substrate itself, as instructions for whoever updates it: a person on your team, or an AI assistant that proposes updates for a person to approve.

- **Record authoring:** Each substrate should include its canonical name, known aliases, key attributes, and relevant relationships, so AI finds the same details in the same place every time.

- **Alias resolution:** Map shorthand, legacy names, and alternate terminology back to a single canonical term, so “CalSTRS” and “California State Teachers Retirement System” are treated as one investor.
- **Relationship validation:** Check that hierarchies, groupings, and relationships between records still agree with each other. If an investor moves under a new parent entity, every record that references it should reflect that change.
- **Structured views:** Keep any comparison tables or structured views that downstream workflows rely on up to date.
- **Source reconciliation:** Periodically compare this file against your systems of record and flag any discrepancies for the owner to resolve.

TIP

Store substrates where your team already works

The tooling for storing and managing your context layer is still evolving, and no single platform has won out just yet. GPs we work with are using tools like Notion, Glean, Guru, or Confluence, or even something simpler: Markdown files sitting in Google Drive or OneDrive/SharePoint. Markdown stored in a code repository like GitHub is another route, with better versioning and permissions than most, but it’s also the hardest for a human to open and actually read.

Every option has tradeoffs: how tightly it controls permissions, how easy it is to update, and how legible it is to humans, in addition to AI. The right choice has less to do with which tool is “best” and more to do with where your team already works.

Keeping the context layer apart from where the work happens forces a choice: either people change where they do their own work, or you build a reconciliation process so the context AI reads doesn’t drift from what your team is actually using. Many horizontal B2B software companies are building against this problem now, so a better answer may arrive in one to two years and make a custom-built solution obsolete.

For context likely to be reused across workflows—such as company terminology or formatting guidance—it can be helpful to store knowledge in smaller, discrete files (what we internally call "atoms"), rather than one giant document. Atoms organize your information so AI can use large volumes of it efficiently. AI can only hold so much in view at once (its "context window"), so instead of having it read the whole encyclopedia every time, let it navigate quickly to the entry that solves the current task.

Turn repeatable instructions into skills

Not all knowledge is reference material. Some of it is instructions, and a skill is how you store them: another kind of context, saved as a reusable prompt along with everything it needs. For a one-off task, you can give AI instructions directly. But if you find yourself repeatedly giving it the same instructions, examples, and supporting context, make them usable as a skill.

A skill packages the instructions and resources that teach AI how to perform a particular task. It can include the steps AI should follow, relevant context files, examples of good output, templates, and other resources it needs to do the job well. For example, an investor tearsheet skill might tell AI which information belongs in the tearsheet, how to structure it, which context files to reference, how your firm defines key metrics, and what a good finished tearsheet looks like. Skills were first developed by Anthropic (makers of Claude), but have now been adopted by every major LLM provider.

This is also where the workflow brief you created in Step 2 and the institutional knowledge you surfaced in Step 3 start to come together. You've defined the outcome, documented how your firm approaches the work, and captured the context AI needs. A skill packages those pieces into something AI can use again and again. ← See Step 2, p. 11 ← See Step 3, p. 13

Skills work at two levels: 1) a skill can be a complete workflow on its own, like a tearsheet creation skill; 2) a skill can be a building block in a larger workflow, where you tell AI to run a skill as one step in a chain of actions.

TRY IT

Turn a standard operating procedure (SOP) into a reusable skill

Start with a task your team performs repeatedly and already knows how to do well. Gather the workflow brief you created in Step 2, the knowledge you surfaced in Step 3, and a few examples of strong past work. ← See Step 2, p. 11 ← See Step 3, p. 13

A useful skill should capture:

- [ ] **When to use it:** What request or situation should trigger the skill?
- [ ] **What to do:** Lay out the process step by step, including the judgment calls an experienced employee makes along the way.
- [ ] **What context to use:** Include or point to relevant context files, examples, templates, and other resources.
- [ ] **What good looks like:** Give AI examples of strong outputs and criteria for evaluating its work.
- [ ] **What to check:** Tell AI what to verify before considering the task complete.

In ChatGPT: Open Plugins → Skills → Create → Create with chat, or simply ask ChatGPT to create a skill for the workflow. Give it the materials above and describe what you want the skill to accomplish; ChatGPT's skill creator can help turn them into a reusable skill. You can test and refine the skill, then share it with others in your workspace if appropriate.

In Claude: Go to Customize → Skills → + → Create skill. You can create the skill there or upload one you've already built. At minimum, a Claude skill contains a SKILL.md file describing what the skill does, when Claude should use it, and the instructions it should follow; you can add supporting files, examples, or scripts as needed. On supported Claude plans, you can also record yourself performing a workflow in Claude for Mac and have Claude propose a skill based on what it observes.

Remember: Start simple. Your first skill doesn't need to automate an entire business process. Once you have a working skill, test it against several real examples, see where it gets the process or judgment wrong, and update the instructions and supporting context accordingly.

WORKSHEET

Build your first substrate

Substrate name _____ **Substrate owner** _____

Purpose and scope: What knowledge or questions does this substrate cover, and what's out of scope?

_____

_____

KNOWLEDGE SOURCES

Knowledge	Canonical source	Internal owner	Status

When two sources disagree, defer to:

_____

_____

Maintenance: When should the substrate be updated, and what events should trigger a review?

_____

_____

5

STEP 5 OF 6

Give AI access to data and tools

So far, you've focused on what AI needs to know to perform the workflow. The next question is what it needs to access to carry out the task.

What AI needs to access depends on what you're asking it to do.

Some workflows only need information. An AI workflow preparing an investor tearsheet might need to retrieve contact information from your CRM, fund performance from your accounting system, and recent correspondence from email.

Other workflows also need to take actions in those systems. An investor relations workflow might draft an email in Outlook, update a CRM record, or create a follow-up task. Eventually, you might even allow it to send certain communications directly.

WHAT YOU'LL BUILD

Mapping the data and tools your workflow needs

In this section you'll break your workflow into steps and identify what information AI needs to retrieve and what, if anything, it needs to do for each one. Here's an example mapping:

Workflow step	Information AI needs	Source	Actions AI takes
Understand the investor	Contact and relationship history	CRM	None
Check recent activity	Recent correspondence	Email	None
Check current performance	Fund metrics	Fund accounting	None
Prepare a response	Firm voice and IR guidance	Skill/context	**Draft email**
Record the interaction	—	CRM	**Update investor record**
Follow up	—	Task manager	**Create task**

This exercise also helps separate the knowledge you captured in Step 4 from information that needs to be retrieved at runtime. ← See Step 4, p. 16 Your firm's communication guidelines might belong in a skill; an investor's latest activity belongs in the system where that activity is recorded.

Start manually when you can

For an early prototype, uploading files or copying in the relevant information is often enough. If you want AI to analyze quarterly reports, reconcile two spreadsheets, or draft a response based on an investor's history, give it representative data and see whether it can reliably produce the result you defined in Step 2. ← See Step 2, p. 9

You can take the same approach with actions. Have AI draft the email rather than send it, or propose a CRM update rather than make the change itself.

This gives you a chance to validate the workflow before granting access to production systems or investing in integrations. Once manually moving information or carrying out AI's proposed actions becomes the bottleneck, it may be time to connect the underlying tools.

Connect the systems the workflow needs

Begin by checking whether your AI tool already offers a connection to the system you need. In ChatGPT, these integrations are called apps; in Claude, they're called connectors.

Depending on the integration, a connection can give AI access to information, tools it can use, or both. One connection might let AI search your email. Another might let it retrieve a CRM record and update it. The important question isn't simply whether a connection exists, but what capabilities it exposes. When evaluating a connection, check:

- [ ] **What can AI read?** Can it search across the system, retrieve individual records, or access only certain types of data?
- [ ] **What can AI do?** Can it create or modify records, draft or send communications, or trigger other actions?
- [ ] **What permissions does it inherit?** Does AI have access only to information the individual user can see, or does the connection use a broader service account?
- [ ] **What requires confirmation?** Which actions can AI take directly, and which require a person to approve them first?

Start with the minimum access your workflow needs. If AI only needs to retrieve investor information, there's no reason to give it permission to modify investor records.

If ChatGPT or Claude doesn't already offer the connection you need in their marketplace, check whether the software vendor provides an MCP server.

DEFINITION

MCP, or Model Context Protocol, is an open standard that allows AI applications to connect to external systems. An MCP server can expose both information AI can retrieve and tools AI can use. For example, a CRM's MCP server might let AI search for an investor, retrieve the investor's activity history, and update a record.

If a vendor provides an MCP server but isn't in the LLM's marketplace yet, it will typically give you a URL and instructions for connecting it. You add that server to your AI tool as a custom connection, authenticate your account, and AI can then use the capabilities the vendor has exposed. Advanced teams can also consider building a custom MCP server on a vendor's API.

TRY IT

Connect a system to ChatGPT or Claude

- [ ] **Look for an existing connection.**

  Search ChatGPT's Apps Directory or Claude's Connectors Directory. If the system is available, connect your account and check which data and actions the integration provides.

- [ ] **If it isn't there, see whether the software vendor offers an MCP server.**

  Search the vendor's documentation for "MCP." If it provides a remote MCP server, follow its setup instructions.

  Claude: Add the server under Customize → Connectors → + → Add custom connector, then authenticate.

  ChatGPT: On supported workspace plans, an admin or an authorized developer can add, test, and publish a custom MCP-powered app to the workspace.

- [ ] **Test the capabilities you actually need.**

  Don't stop once the connection works. Ask AI to retrieve the information your workflow requires and, if applicable, test the actions it needs to take. Make sure the connection exposes the right capabilities and respects the permissions you expect.

Start with read access where possible. Add actions such as updating a CRM record or sending a message only when the workflow requires them.

The built-in connections in ChatGPT or Claude may be enough for your first workflows. If you're doing this at enterprise scale, you might want to consider a tool gateway such as Runlayer or MintMCP. A gateway like this adds a governance layer between AI and the systems it can access, giving security teams more control and visibility into which tools agents are using and giving employees a library of approved integrations they can reuse across workflows. You probably don't need one to get started, but a gateway becomes more useful as the number of systems, workflows, and people building with AI grows.

Build a custom integration only when you need one

If the system you need offers neither a ready-made connection nor an MCP server, decide whether the workflow is valuable enough to justify building the connection yourself.

A developer can build an MCP server that exposes the specific information and tools your workflow needs. If you're building a dedicated application, you may instead integrate directly with the system's API, giving you tighter control over what information the application retrieves, how it processes that information, and what actions it can take.

You don't need to make this architecture decision upfront. A workflow might start with uploaded files and human-executed actions, move to existing connections as it proves useful, and eventually warrant a custom integration as it becomes more important to the business.

TRY IT

Connect Juniper Square to your AI tool

Juniper Square offers an MCP server that lets supported AI tools retrieve data from Juniper Square and use it as context for your work. That means you can ask AI questions about your firm's data without exporting and uploading it each time.

To get started:

- [ ] **Search for Juniper Square in your AI tool's marketplace.** If your LLM doesn't have our MCP in its marketplace yet, reach out to your account team and we can help you get set up.
- [ ] **Authenticate with your Juniper Square Enterprise credentials.**
- [ ] **Test the connection with a real workflow.** Start with a question that requires current Juniper Square data rather than something you could answer from a static document. For example, ask AI to create an investor-specific tearsheet that brings together the investor's profile, commitment (including called capital and unfunded commitment), fund exposure, relationship history, and fundraising context.

WORKSHEET

Map the data and tools your workflow needs

Now it's your turn. Using the information you've gathered in this step, fill out the worksheet below. If you need help, refer back to the example on p. 24.

Workflow step	Information AI needs	Source	Actions AI takes

STEP 6 OF 6

6 Decide how to implement the workflow

Once you've defined the workflow and given AI the knowledge, data, and tools it needs, decide how to put it into practice.

There are three primary ways to build an agentic workflow: a skill, an agent, or a dedicated application. Start with the lightest-weight option that can reliably do the job. If a skill is enough, stop there. Build an agent or a dedicated application only when the workflow genuinely requires greater automation, orchestration, scale, or control. If you want colleagues to use a skill or agent you've built, you can package it as a plugin (see the tip on p. 34).

As you choose, consider how the workflow will be used, how independently it needs to run, and how it treats data:

- A workflow that someone invokes as needed inside an AI client may only need a skill.
- If it needs to run automatically or make decisions across multiple steps, you may need an agent.
- If you want colleagues to use a skill or agent you've built from inside their own AI client, without a separate application, you can package it as a plugin.
- And if it needs its own interface, stores data that doesn't already live in another system of record, or requires sharing the tool with others to be useful, a dedicated application may make more sense.

Note that these approaches aren't mutually exclusive. An application might call an agent behind the scenes, while an agent might use several skills and tools to complete its work.

Option 1: Use a skill for a repeatable task

A skill packages the instructions, context, and resources AI needs to perform a particular task consistently. It works well when a person is still initiating the work and interacting directly with AI. An IR team member, for example, might invoke a skill to build an investor tearsheet, prepare for an LP meeting, or draft a response according to the firm's guidelines.

If you've already created a skill in Step 4 and it reliably does what you need, you may be done. ← See Step 4, p. 20 You don't need to turn every AI-enabled workflow into something more sophisticated.

Option 2: Use an agent to carry work across multiple steps

An agent makes sense when you want AI to take responsibility for a larger portion of the workflow. Rather than waiting for an employee to initiate each task and tell it what to do next, an agent can respond to a trigger or schedule, work toward a goal across multiple steps, decide which information and tools it needs along the way, and take actions on your behalf.

An investor relations agent, for example, might monitor an inbox, classify an incoming question, retrieve the relevant investor and fund data, apply your investor-response skill, create a draft in Outlook, and route it to the IR team for approval.

TRY IT

When to use an agent

Before you build an agent, make sure your workflow has:

- [ ] **A clear trigger:** An investor sends an email, a new document arrives, or a reporting deadline approaches.
- [ ] **A defined goal:** Respond to the inquiry, review the document, or prepare the report.
- [ ] **Tools it needs to use:** Your CRM, investor portal, email, accounting system, or other systems.
- [ ] **Decisions along the way:** The right next step depends on what the agent finds.
- [ ] **Actions it can take:** Draft an email, update a record, create a report, or route something for review.
- [ ] **Clear boundaries:** You know what the agent can do independently and when it should stop or ask for approval.

In ChatGPT: Agents can combine instructions and skills with tools, MCP connections, code execution, and other capabilities. Depending on the use case, you might package an agent as a plugin that employees can access directly in ChatGPT, or build it into a separate application (we'll talk more about these below). OpenAI offers both managed agent infrastructure and software development kits (SDKs) for teams that want more control over how agents run. See p. 34 →

In Claude: Agents can be combined with tools and MCP connections to carry out multi-step work. You can make these capabilities available inside Claude or use Anthropic's developer tooling to build an agent into a more customized application.

For more complex agents, orchestration platforms such as Workato or Microsoft Copilot Studio can provide infrastructure for triggers, integrations, error handling and retries, human approvals, and coordinating multiple agents within a workflow.

TIP

Package capabilities as a plugin

If you want to make what you've built available to other employees inside their existing AI client, you may be able to package it as a plugin rather than building a separate application. In Claude, for example, a plugin can bundle skills, connectors, agents, and other resources into something that can be shared across a team.

This can be a useful middle ground when the workflow works well inside an AI client and doesn't need its own interface. Employees get a reusable, preconfigured way to access the capabilities you've built without needing a separate web application.

Option 3: Build a dedicated application when the workflow needs its own product

Some workflows need more than a skill or agent running inside ChatGPT or Claude. A dedicated application is a piece of custom software that incorporates AI. It may make sense when the workflow needs its own interface, will be used broadly across the firm, needs to maintain its own data or application state, or requires tighter control over the user experience and underlying infrastructure. Building one well takes more than understanding AI. It also takes an understanding of software design principles, including security best practices, so the application doesn't become a weak point in your firm's cybersecurity posture.

A web application can also make a workflow easier to distribute. Instead of asking coworkers to configure or invoke something inside an AI client, you can give them a link. If ease of access is important to adoption, that alone may be a reason to package an otherwise simple workflow as an application.

TRY IT

Should you build a dedicated application?

Most AI applications have a few basic components. Check each question you can answer yes to. Anything left unchecked is work to plan before you build.

- [ ] **Do you know where people will use it?** This could be as simple as a form or chat interface, or a more purpose-built experience like an investor dashboard.
- [ ] **Have you chosen a model?** It reasons about the task, generates outputs, or decides what to do next.
- [ ] **Are the instructions and context written down?** These are the skills, prompts, context files, and examples that teach the model how your firm performs the work.
- [ ] **Can it reach the tools and data it needs?** It may need to read from or act in systems such as your CRM, fund accounting platform, investor portal, or email.
- [ ] **Have you defined the application logic?** Decide when the workflow runs, what steps happen in what order, and what to do when something goes wrong.
- [ ] **Do you know whether it needs its own database?** It does if it creates or maintains information that doesn't already live in another system of record.
- [ ] **Have you decided who can use it?** Set who can access the application and which data and actions each person can use.
- [ ] **Will you know when it fails?** You need a way to see what the application did, identify failures, and measure whether its outputs continue to meet your quality bar.

Choose how you'll access AI models

For a straightforward application, the simplest option is often to work directly with a model provider such as OpenAI or Anthropic. But model capabilities, pricing, and availability are changing quickly. If you're building something you expect to use and expand over time, consider how easily you'll be able to evaluate or switch models as better options emerge.

A managed model platform or gateway, such as AWS Bedrock or OpenRouter, can give your applications access to models from multiple providers through a common layer. This makes it easier to choose different models based on capability, cost, or availability without rebuilding the application around a new provider.

Setting this up requires some technical work. You'll need to decide which models and providers your applications can use, configure any routing or data-location requirements, and control which users and applications can invoke them. Depending on the complexity of your setup, you can configure this through the platform itself or programmatically through an API or SDK.

Consider model cost and control

Relying on a hosted model provider means the provider ultimately controls which models remain available and how they're priced. As new models are released, older models may be retired or replaced with models that have different capabilities and economics.

For applications where cost, control, or model availability matters, open-weight models provide another option. Models such as Meta's Llama, Mistral, DeepSeek, or Google's Gemma can be run on infrastructure you control or through a managed platform. That gives you more control over how and where the model runs, but also means someone needs to handle hosting, compute for inference, security, and ongoing operations. For most firms, using a managed platform such as AWS Bedrock, Databricks, Azure AI, or Google Vertex AI will be more practical than operating this infrastructure themselves.

There are also emerging models designed specifically to make narrow, repeatable automations cheaper. Jev, from TypeSafe AI, takes a different approach from a general-purpose LLM: rather than generating free-form text, it returns probabilities for a predefined set of outcomes. For example, instead of asking an LLM to determine whether a lead already exists in your CRM and generate a written answer, you could define `true` and `false` as the possible outcomes and use Jev's probabilities to determine what the workflow should do next. That makes this kind of model potentially useful for high-volume classification and decision steps inside larger workflows, where you don't need an LLM to generate or reason through a long response.

How to build it today

AI tools have made it much easier to build a first version of an application without a large engineering team. App-building platforms such as Lovable or Replit can take you from a natural-language description to a working web application, including the interface and much of the underlying application logic. They can be a good place to start when you want to prototype an internal tool quickly.

Coding agents such as OpenAI Codex or Claude Code give you more direct control over the underlying code. They can help you build interfaces and application logic, connect APIs, write tests, and work with an existing codebase from natural-language instructions.

The line between these categories is increasingly blurry, and either approach can get surprisingly far. The more important distinction is between prototyping something and operating it as a production system. If an application will handle sensitive data, take actions in production systems, or become part of a critical business process, you'll want someone with the technical expertise to review the architecture and code, configure security and permissions, test failure cases, and own the application in production.

WORKSHEET

Decide how to implement your workflow

Start with the lightest-weight option that can reliably do the job. If a skill is enough, stop there.

Your workflow

CHOOSE YOUR OPTION

How will the workflow be used?

- [ ] Someone invokes it as needed inside an AI client — **Skill**
- [ ] Colleagues use it from inside their own AI client — **Plugin**
- [ ] It needs its own interface — **Application**

How independently does it need to run?

- [ ] A person initiates the work — **Skill**
- [ ] It runs automatically or makes decisions across multiple steps — **Agent**

How does it treat data?

- [ ] It uses data from your existing systems — Any option
- [ ] It stores data that doesn't already live in another system of record — **Application**

Your starting option

- [ ] Skill p. 32
- [ ] Agent p. 32
- [ ] Plugin p. 34
- [ ] Dedicated application p. 34

IF YOU NEED AN AGENT

A clear trigger

A defined goal

Tools it needs to use

Decisions along the way

Actions it can take

Clear boundaries

BEFORE IT GOES INTO PRODUCTION

Will the workflow:

- [ ] Handle sensitive data?
- [ ] Take actions in production systems?
- [ ] Become part of a critical business process?

If you checked any, who will review the architecture and code?

What this playbook looks like in action

Atlas Real Estate Partners has found more than one way to embed AI into its workflows, and has done so without a dedicated AI technologist in-house. Instead, the team's AI transformation has grown out of genuine curiosity about how AI tools work and frustration with inefficient, high-toil workflows that became impossible to ignore.

One workflow the firm has transformed is drafting responses to inbound questions for its investor relations team. Each day, a purpose-built AI agent scans a team-wide email alias for new inquiries and drafts responses directly in Microsoft Outlook for the IR team to review. The agent can also handle routine administrative tasks for the IR team, such as updating a contact's preferences in the investor portal or retrieving documents from the portal and attaching them to a draft email.

The knowledge behind those responses came largely from material Atlas already had and continues to enhance. Years earlier, Zach Chilelli, VP of Investor Relations, had written SOPs in Microsoft Word to train new employees on common investor questions. He had also recorded videos explaining more complex workflows and processes. He used that material to teach agents how Atlas handles different questions and to make their responses more consistent with the firm's brand and best practices.

The workflow also has a clear boundary around what AI can do. The agent drafts the response, but doesn't send it. Instead, its emails sit in the drafts folder until a member of the IR team reviews them, makes any necessary changes, and sends the final response.

The result brings together many of the practices in this playbook: Atlas started with a repetitive workflow worth improving, turned existing institutional knowledge into instructions AI could use, connected the agent to the systems where the work happens, and kept a team member responsible for the final output.

The example also shows why writing it down matters. The SOPs and training materials Atlas had already created for people became the foundation for teaching AI how the firm works.

> "It's been truly transformational for the IR department at Atlas."
>
> **Zach Chilelli**, VP of Investor Relations, Atlas Real Estate Partners

PART 2

Top workflows for AI transformation

Four areas where GPs are seeing real-life benefit from AI in fund operations

Picking the right workflows to focus your efforts on is perhaps the most important step to get right. Teams often start with an overly ambitious goal and then realize deep into the project that AI can't deliver the accuracy rates they'd need to actually trust it in production.

We've surveyed and worked with hundreds of our customers, and clear themes are emerging in where AI is changing workflows. The four categories below are the most common, and the workflows in them are ones we know can work. This list isn't exhaustive, but it can help you pattern-match with the areas of greatest need at your firm.

1 **Keep investor and fund data consistent and audit-ready** p. 41

2 **Improve the efficiency of investor relations** p. 42

3 **Empower finance and operations teams to make faster decisions** p. 43

4 **Bring the team together** p. 44

1 Keep investor and fund data consistent and audit-ready

Investor, operational, and fund data often live in different systems, and keeping them aligned is manual work that AI can streamline, helping IR and finance teams work hand in hand.

EXAMPLES

- **Reconcile data across systems:** Compare investor and fund data across systems, identify records that don't match, and investigate the likely source of the discrepancy. Instead of simply syncing fields, AI can use surrounding context to flag differences that need attention and help determine which information is authoritative.
- **Build management company financial views on demand:** Give AI access to expenses, travel, management fees, and other operating data, then use it to build forecasts or custom analyses as needed without maintaining every view in a separate FP&A tool.
- **Reconcile and audit tie-outs:** You can also reconcile your general ledger against wherever your transactions live (including Juniper Square). AI can automatically pull that data and surface the specific line items that don't tie out, so your team knows the potential threads to chase.
- **Analyze documents at scale:** Extract and compare terms across LPAs, subscription documents, and underlying GP reports, flagging inconsistencies or unusual provisions for review. For higher-stakes workflows, require the model to cite its source, distinguish between fields it can extract confidently and those that need human review, and test the output with evals (see p. 51).

FROM THE FIELD: RIBBIT

The team at Ribbit delivers more than 50 different financial statements to its investors each quarter. Using AI helps ensure a high-quality product, catching formatting and calculation inconsistencies that could be missed when reviewing them by hand, says Ben Paull, CFO of Ribbit. "It's getting the little things right that creates trust in the relationship."

2 Improve the efficiency of investor relations

Free your IR team to spend more time on the relationship side of investor relations by having AI take over the manual work.

EXAMPLES

- **Create investor tearsheets:** Generate a personalized tearsheet built to your firm's preferences. With AI, you can choose which metrics, interaction history, and portfolio details appear for institutional LPs versus family offices.
- **Automate email ticketing:** Draft responses to investor questions automatically, grounded in your own internal data rather than a generic template.
- **Screen and prioritize prospects:** Evaluate a list of prospects, such as conference attendees, against your CRM, past communications, investment activity, and other relevant signals to identify the relationships most worth pursuing and surface the context your team needs to act on them.
- **Look up distributions and capital activity:** Speed up the answers to questions about positions, transaction events, and distribution notices, such as what a specific investor received or what's unfunded.
- **Build a complete relationship history:** Reconstruct the history of a prospect or investor across multiple people and conversations by giving AI access to firm-wide communications.

FROM THE FIELD: ATLAS REAL ESTATE PARTNERS

Zach Chilelli, VP of Investor Relations at Atlas Real Estate Partners, built a CRM enrichment tool that uses AI to monitor LinkedIn updates, SEC 13F filings, and industry publications to keep more than 4,000 prospect and investor contacts up to date. Instead of someone manually checking whether a filing or a headline means anything for a given relationship, the AI reads and watches for a trigger to be met, such as a contact receiving a promotion or starting a new venture. Then it sends an alert when it's worth reaching out.

3 Empower finance and operations teams to make faster decisions

Give finance and operations teams direct access to the information they need to spot issues and more quickly resolve them.

EXAMPLES

- **Investigate financial variances:** Ask AI to review financial data for anomalies, drill into the transactions behind a variance, and surface the records most likely to explain it. Finance teams can investigate an unexpected number from the top-line result down to the underlying transactions without manually tracing it across systems.
- **Analyze performance across funds:** Roll up performance across funds and ask questions across the resulting data, for example, which investors have increased commitments over time.
- **Monitor your fundraising pipeline and data room:** Code a custom dashboard that's integrated with the data sources you care about. For example, an IR dashboard might merge data from your CRM and Outlook and update automatically as investors email you.
- **Track your org chart and entity structure:** Map the legal entity and org structure for each fund, and keep it up to date as things change.
- **Customize business intelligence and reporting:** Replace your business intelligence dashboard with a generative interface. Tell AI the view you want, and AI will build it on demand instead of you maintaining a fixed set of reports.

FROM THE FIELD: ATLAS REAL ESTATE PARTNERS

The team at Atlas Real Estate Partners used AI to write the code for a capital-raising dashboard that tracks five active raises. AI reads inbound and outbound emails and pulls out the stage of each conversation. Anyone on the team can see who's a hot prospect, who's still waiting on a reply, and who declined and why, without any manual updates.

4 Bring the team together

When AI is embedded in a firm's culture, it can increase access to information and make the team more unified and aligned.

EXAMPLES

- **Tailor information to how people absorb it:** Turn industry articles, people updates, and CRM notes into something people will actually consume. That might be an audio recap for their commute to the office, a weekly summary, or a video clip that gets played during an all-hands meeting.
- **Turn institutional knowledge into reusable workflows:** Automate common workflows by recording how someone does a specialized task, such as valuing a specific type of interest or walking through an edge-case transfer, and then teach an agent how to do that task consistently.
- **Track what's happening in the office:** Tell an agent to look at systems with valuable data that informs daily events, such as a shared calendar, Slack channels, and a booking system. From that, it can build and maintain a living reference of what's happening in the office, so employees don't have to ask the office manager which portfolio companies or LPs are visiting in a given week.
- **Build organizational memory automatically:** Point AI at every Zoom or Teams transcript across the firm and ask it to build a record of decisions, commitments, processes, and institutional knowledge that employees—and eventually AI workflows—can draw on.

FROM THE FIELD: RIBBIT

The team at Ribbit produces a weekly internal-facing AI-generated podcast to share the information the team finds most insightful. Each episode draws on articles, social media posts, and other material the team has flagged, so most of the team can catch up quickly.

CASE STUDY

How Ribbit built a culture around AI

AI is built into how Ribbit operates day-to-day. Each week, team members share with the group the projects they've built with AI that are making their work better or more efficient. They share how they built it and what it's let them do, which speeds up everyone's learning curve. Meaningful tools, like an AI-generated weekly podcast, come from these demos and have a firm-wide impact.

CFO Ben Paull's finance team is where AI faces its hardest tests, since finance is the one function at the firm with no room for error. You can't get a financial statement wrong or send an investor a report with a mistake. Connecting fund and investor data directly into the AI tools used by his team improved the consistency and output of deliverables. Paull and his team can now build their own reports on demand, review financial statements for consistency across dozens of filings at once, and generate management company financials without maintaining a separate FP&A system.

> "When everyone knows that the whole group is working together on this, they really dig in."
>
> **Ben Paull**, CFO, Ribbit

PART 3

A primer for executives managing AI initiatives

Building a useful workflow is one part of adopting AI. As AI takes on a larger role in day-to-day work, you'll also need to think about how you encourage productive usage, maintain the quality of the work your firm produces, evaluate what's working, and manage the risks that come with greater access and autonomy.

You don't need to solve all of these questions before you get started. But as your use of AI grows, a few firm-wide practices can help you scale what works without losing the judgment, quality, and control your business depends on.

Give AI work clear owners

AI projects work best when someone is responsible for the outcome, not just for getting the technology working.

In practice, most projects need three roles:

- **Domain expert:** Someone who understands the workflow and can tell whether AI is doing the work correctly.
- **Solution architect:** Someone who can translate the business requirements into an appropriate systems and knowledge architecture.
- **Builder:** Someone who can take the specified architecture and build a high-quality workflow.

The solution architect is the role firms are most likely to miss. There are plenty of builders now, but far fewer people who combine domain and AI expertise well enough to make the right architectural choices.

These don't need to be three different people. For a simple skill, an AI-savvy employee who knows the workflow may play all three roles. More complex applications may require an engineer or AI specialist to build the system while a domain expert owns the requirements and evaluates its work.

These three roles parallel how the best startups work: domain expertise, great product management, and great engineering.

Executive sponsorship may sit outside these three, but it matters just as much: ownership of the outcome stays with the business. The person building an AI workflow can help encode how the work should be done, but someone who understands the underlying business goals needs to be able to say whether it's actually delivering value.

Get the whole team to adopt

Choosing tools and getting started with AI is relatively easy. Getting the whole team to incorporate AI into how they work is harder. Three things can help:

1. **A visible executive champion:** When a firm's leadership uses AI tools themselves, publicly and regularly, it gives the rest of the organization permission to experiment. People are more willing to try new ways of working when they've watched someone senior do it first and work through the mistakes.
2. **A small working group:** Start with three to five people who are already comfortable with new technology. They can build the first workflows, share what they learn, and become internal resources for other employees.
3. **Hackathons or demos:** Give people low-stakes opportunities to test ideas and show what they've built. Some of the most useful applications emerge when employees start experimenting with problems they encounter every day.

A note on executive champions: there's a real difference between a leader who's genuinely working through the complexity and one who's just performing enthusiasm for it. The most useful kind of champion learns what it actually takes to do this work well, including the judgment calls, resourcing, and the places where AI isn't the answer yet, and is then honest with their team about what they've learned. The unhelpful kind just pushes AI on everyone and ends up with a firm full of rushed work dressed up as progress.

EXAMPLE

At Juniper Square, CEO Alex Robinson acts as AI's executive champion. He holds a weekly meeting to discuss AI transformation, prioritizing alignment and ownership. By getting stakeholders in a room together and using an AI Chief of Staff to record decisions and action items, we reduce the coordination tax and accelerate AI progress.

Standardize the tools your team uses

As AI usage spreads, it's worth giving employees a common set of tools rather than leaving everyone to assemble their own stack of consumer AI accounts.

Start by distinguishing the AI client your employees work in from the model infrastructure you might use to power custom applications. For employees, the client is where they'll analyze data, create content, work with company information, and increasingly build lightweight workflows and tools.

Integration with your existing workplace software matters. A firm already deeply invested in Microsoft products, for example, may naturally consider Microsoft Copilot. But also look at what the AI client lets employees do beyond prompting for answers. Can they create reusable skills? Connect AI to other systems? Build lightweight applications? Work with agents? These capabilities increasingly determine how far employees can take AI without needing a dedicated engineering team.

In our experience, firms getting the most out of AI tend to give employees access to general-purpose AI clients such as ChatGPT or Claude, where many of the newest capabilities for connecting systems and building AI-enabled workflows are appearing first.

TIP

Choosing an AI plan for your firm

For smaller firms without specialized IT or compliance requirements, a business or team plan is often a reasonable starting point. These plans are typically priced per seat and include centralized administration, authentication, and protections for company data.

Enterprise plans become more relevant when you need capabilities such as automated user provisioning, granular access controls, custom retention policies, audit logs, or additional compliance and security controls.

Before choosing a tier, have whoever owns IT or compliance identify the requirements that actually matter to your firm, then compare them against the providers' current offerings. Pay particular attention to how the provider handles business data, what information may be retained, and how your data is used.

Both OpenAI and Anthropic publish documentation specifically covering how data from their business products is handled. Include those policies in your security or compliance review rather than relying on assumptions about how consumer AI products work.

Encourage the right AI usage

Once people have access, resist the temptation to make raw usage the goal.

Look at your usage dashboards to find your strongest natural AI users and understand what they're doing, not to rank employees against one another. The more you measure and reward raw usage, the more likely you are to get usage for its own sake.

Instead, look at what people are accomplishing. High usage may indicate that someone is inefficiently using an expensive model for a simple task, but it may also mean they're building a valuable application or automating a workflow that previously took hours. The conversation should start with what they're doing and whether it's useful.

Usage data can also surface training opportunities. If someone consistently uses a slower or more expensive model for straightforward work, show them when a faster model will do the job. If they're reaching usage or spend limits because they've found a valuable use case, that may be a reason to give them more capacity rather than less.

TIP

Match the model to the work

Most major AI providers now offer a range of models that trade off capability, speed, and cost.

As a general rule, use a fast general-purpose model for everyday work and reserve the provider's most capable reasoning models for harder problems where the additional intelligence materially improves the result. Very lightweight models can handle simple retrieval, classification, or repetitive tasks.

Revisit your guidance as new models ship. A task that requires the most capable model today may be handled well by a faster, cheaper model weeks from now.

Avoid AI slop

AI is genuinely useful for drafting investment memos, LP letters, and the other written work that fills a GP's week. But making writing easier to produce also makes it easier to produce work that is generic, verbose, or disconnected from the judgment your firm is being paid to provide.

A few practices can help:

Before you write, do the thinking yourself. The value of a memo is in the human judgment behind the words, including what you focus on and the risks you flag. Use AI once you know what you want to say and can provide it with your genuine thinking, rather than outsourcing the underlying judgment.

Make the human responsible for the output. Whoever generates a piece of writing with AI should still own what it says. When people are paying for your judgment and critical thinking, make sure the value your team adds doesn't get lost.

Only use as much space as necessary. AI makes it easy to generate more. That doesn't mean a memo should be longer than it would have been if you wrote it yourself. Keep a document as short or as long as the substance requires.

Proofread the output. Before sending anything, read it closely enough to make sure it says what you mean in both substance and language. If a sentence doesn't sound like something you would write, remove it or rewrite it.

Turn your brand guidelines into a plugin. For firms that have developed brand guidelines, including voice and tone attributes, turn these materials into a plugin that everyone can use to ensure that anything AI writes aligns with your existing style guide.

Evaluate whether AI is working

Once AI is producing real work, evaluate the output against the success criteria you defined when you built the workflow. This aspect of building is known in the AI world as "evals," and is perhaps the most important part of understanding whether what you've built will be good enough to be used in practice.

There are several useful ways to do this, and many firms will use a combination:

Human review: Have someone who understands the work regularly check AI's output for factual errors, hallucinations, drift, and failures of judgment. Their feedback should flow back into the workflow so its instructions, context, or implementation can improve over time.

AI-graded review: Use a second AI model to evaluate the output against explicit criteria, such as factual accuracy, completeness, adherence to instructions, or the firm's standards. This can scale further than manual review and can be useful as an initial quality check.

Track human changes: Save AI's original output alongside the version a person actually uses. The difference between the two can show you where the workflow needs to improve. If someone consistently changes the same fields in an AI-generated one-pager, for example, that's useful feedback about what the original workflow is getting wrong.

How much review you need depends on the consequences of an error. A brainstorming tool can tolerate mistakes that an investor-facing financial workflow cannot.

TRY IT

Decide how much autonomy to give AI

You don't need to make an entire workflow manual or autonomous. Decide step by step what AI should be allowed to do.

- **Recommend:** AI proposes what should happen; a person decides.
- **Act with approval:** AI prepares the action, but a person approves it before it happens.
- **Act autonomously:** AI takes the action without prior approval.
- **Escalate:** AI stops and asks for help when it encounters a situation outside its instructions or confidence threshold.

An investor relations workflow, for example, might retrieve investor information automatically, draft a response automatically, require human approval before sending the email, and escalate if two systems disagree about a financial figure.

Start with more oversight for consequential actions. As you see how a workflow performs across real cases, you can give it more autonomy where the evidence supports it.

Add guardrails as AI takes on more responsibility

The more access and autonomy you give an AI workflow, the more important it becomes to limit the consequences when something goes wrong.

Limit what each workflow can access

A workflow rarely needs access to your full system of record. Give it only the information and tools it needs to perform its job.

A reconciliation workflow, for example, might need read access to the general ledger and fund administration platform but no access to investor communications and no permission to modify the CRM.

The same principle applies to actions. If a workflow only needs to draft an email, don't give it permission to send one.

Limit how much AI can change at once

Even when AI has permission to take an action, you can limit its blast radius.

A workflow that updates records might be capped at a certain number of changes within a given period. A workflow that encounters an unusually large transaction or a result outside an expected range might stop and require human approval.

The right limits depend on the workflow. The principle is to make sure one bad input or misunderstood instruction can't create hundreds of errors before anyone notices.

Keep a record of what AI did

When one person performs a workflow manually, they often know what changed and why. Once AI is doing that work across the firm, that context can disappear unless you preserve it.

For consequential workflows, keep a record of the information AI used, the actions it took, and where human approval was involved. That gives you an audit trail when something goes wrong and better information for improving the workflow over time.

You don’t need the same controls around every use of AI. A skill that helps an employee brainstorm questions for a meeting doesn’t need the same oversight as an agent that can update investor records. Add controls in proportion to what the workflow can access, change, and communicate.

CONCLUSION

Your data is the head start

Putting AI to work in your firm starts with individual workflows: choosing work worth improving, defining what good looks like, giving AI the right context and access, and deciding how much of the workflow it should take on. As those workflows multiply, the focus expands to the broader firm: who owns the work, which tools your team uses, how you maintain quality, and how you evaluate and govern what AI is doing.

You don't need to have all of this figured out before you start. The firms making the most progress are building useful workflows now, learning from how they perform, and using those lessons to shape how they adopt AI more broadly.

One advantage private markets firms already have is the context AI needs to do useful work. Public markets run on information that's broadly shared: the same filings, the same prices, the same disclosures available to everyone. Private markets don't work that way. Every firm has its own strategy, relationships, and data, and that's a real part of why this industry has resisted the kind of commoditization public markets have seen.

No AI model can tell you what the market looks like for management fee waivers on a $1 billion PE fund, because that knowledge doesn't live anywhere public. It lives with the people at your firm raising your next fund, who hear how LPs react or talk to the lawyers who see patterns forming across deals. Getting ahead with AI means pulling that knowledge out of your systems of record, operating processes, investment frameworks, and the heads of the people who carry it. Then making it into something AI can actually use.

Building AI tools on a strong foundation

A general AI model is only as good as the context you give it, and a lot of that context, including the data behind how your firm actually runs its funds, already lives inside Juniper Square.

Juniper Square holds the data behind more than 2,300 GPs' fund operations: the contacts, positions, transaction history, permissions, and audit trail that make that data trustworthy. As AI gets more capable, well-structured, reliable data is what gives firms something solid to build on.

We've spent the past year going through an AI transformation ourselves, and we've shared many of the lessons we learned along the way in this guide. If you're figuring out what to build next or how to put your data to work with AI, we're here to help. Reach us at ai-transformation@junipersquare.com.