TRUE NORTH INSTITUTE

# AI Tools for Institutional Investors - a Whitepaper

(June 2026)

## Executive Summary

*The focus of this paper is on how institutional investors ("asset allocators") can presently get the most value out of available AI tools. This does not attempt to address what asset managers (GPs) should be doing with AI, nor what AI tech companies LPs and GPs should be investing in.*

*Our recommendations fall into two phases. Phase 1 (~12 months) deploys best-in-class Specialist AI Tools alongside GenAI LLMs to capture immediate productivity gains. Phase 2 (12–36 months) connects those tools into an agentic architecture — linking specialist platforms, CRM and document systems through orchestration agents and Model Context Protocols to create an agentic operating system. Phase 1 is about adopting capabilities that already exist off the shelf. Phase 2 is about building architectural advantage, where multi-step workflows run autonomously and institutional memory compounds with every DD review and IC decision.*

***For asset allocators, the four core near-term Phase 1 use cases are manager due diligence, portfolio construction, workflow management and data infrastructure.*** *The best AI tools differ across each category. Near-term value is concentrated in manager due diligence and data infrastructure. Portfolio construction remains a hybrid AI + human workflow, while workflow management is evolving through incremental AI enhancements to existing platforms rather than AI-native replacements.*

***In manager due diligence*** *and data room mining, GenAI LLMs such as Anthropic Claude and OpenAI ChatGPT already support much of the first-pass analytical work across legal documents, DDQs, ODD materials, deal tables, team bios, pitchbooks and references. Many institutional investors already use these tools to review legal documents and flag off-market terms or irregularities.*

*However, off-the-shelf LLMs cannot yet ingest and analyse entire data rooms effectively. Specialist AI tools such as* ***MSCI Vantager, Clade*** *and* ***Hebbi****a can ingest full data rooms alongside third-party references, meeting notes, portfolio company financials and internal analyses. These platforms enable effectively 100% document review, versus the historical reality where analysts might only review 20–30% of materials due to time constraints. Asset allocators can define required data fields by asset class, allowing the tools to identify gaps rapidly and generate first-pass investment memoranda for analyst review. This creates the potential for alpha enhancement by enabling investors to go "shallowly deep" across many more managers before going "deeply deep" on a select few. These tools become even more valuable for cross-asset manager comparisons, source-linked analysis and building searchable institutional memory. The ultimate vision for AI-enhanced manager selection is that* ***selection algorithms*** *are developed and evolved from AI finding powerful predictive relationships between manager attributes and consistent long-term alpha generation.*

***In portfolio construction and scenario-based asset allocatio****n, GenAI LLMs are useful for scenario generation, probability framing and IC narrative development, but implementation still relies on legacy end-to-end platforms such as BlackRock Aladdin, MSCI BarraOne and Bloomberg L.P. PORT. A fully AI-driven integrated scenario-to-portfolio solution does not yet exist.*

***In workflow management****, AI currently acts as an enhancement layer on top of legacy systems rather than a replacement. AI is increasingly embedded into workflows such as manager onboarding, GP portal access, compliance checks and capital call reconciliation, but no AI-native platform has yet displaced incumbent workflow infrastructure.*

***In data infrastructure,*** *AI is most valuable when embedded within dedicated systems such as* ***Canoe Intelligence,*** *where the primary challenge is connecting to GP portals, extracting and reconciling data at scale, and feeding clean outputs into downstream systems.*
*The key misconception to reject is that AI will replace manager selection or determine asset allocation decisions autonomously. Its real value today — and likely in future — lies in automating routine analytical tasks: reading DDQs, extracting data from PDFs, reconciling reports and drafting first-pass investment memoranda. This allows investment teams to spend more time applying judgement to manager selection and portfolio construction. The practical rule is straightforward: use GenAI LLMs for one-off analytical tasks; use Specialist AI Tools when scale, ingestion, reconciliation and auditability matter; and rely on AI-augmented legacy platforms for implementation, control and reporting.*

***Culture is critical to successful AI integration.*** *The primary constraint on AI adoption is rarely the technology itself. Institutions at the forefront of AI adoption consistently identify culture as the limiting factor. The fastest-moving organisations typically combine clear top-down AI leadership — including CEO-level mandates, mandatory training and enterprise deployments — with bottom-up experimentation that gives junior professionals explicit permission to redesign workflows around AI tools.*

*The highest-performing organizations are not choosing between humans or machines but continuously optimizing the combination of both. The greater risk is often resistance to AI adoption, not overreliance on it. At the same time, no pure AI output should be considered decision-ready: hallucinations and faulty summaries remain real risks, making human judgment, augmentation, and review essential even when AI produces the first draft.*

*The sequencing lesson is to build early momentum through narrow point solutions and targeted workflow wins — such as co-investment screening tools, first-pass memo drafters or legal document reviewers — before attempting broader platform integration or agentic architectures. Junior* *investment professionals should be encouraged to experiment early and build fluency. Team members who fail to augment their productivity with AI risk displacement by colleagues who do.*

## What AI Tools are Institutional Investors using today?

Few institutional investors are ignoring AI. The question institutional allocators are asking today is where in their workflow AI can contribute meaningfully to alpha generation. Adoption data says we are still in a “wait and see” stage of AI usage by institutional investors. Cerulli's Q2 2025 survey of 200 institutional investors found 12% already incorporate AI into investment office activities and a further 58% are considering implementing the technology.

In April of this year, Capital Allocators held their CIO Summit in Orlando Florida for approximately 120 CIO level investors, including 40 GPs and 80 LPs, with the LPs from about 1/3 E&Fs, 1/3 SFO/MFO, and 1/3 other institutions (SWF, big pension, insurance). In summary, Ted Seides’ meeting concluded that AI in institutional investing today is universally adopted but unevenly implemented, delivering clear gains in productivity and workflow efficiency while leaving core investment judgment, organisational processes, and risk frameworks largely unchanged.

The myth around AI in investing is that it will replace manager selection or set asset allocations. It will not. The more useful framing, articulated by several of the practitioners building in this space, is that AI will automate “grunt work,” while investment professionals exercise judgement, conviction and the final decision on asset allocation and manager selection. In practice, this means compressing the time allocator teams spend reading DDQs, reconciling quarterly reports, extracting data from unstructured PDFs, drafting manager updates and redeploying that capacity to the activities that actually drive returns: manager selection, portfolio construction and macro positioning.

Practitioner adoption is widespread but uneven. Allocators are using AI for drafting investment memos, manager evaluation, research synthesis, note-taking and legal document reviews, but few feel they are moving fast enough. The consistent message is that AI remains a productivity tool, not a full replacement for the existing investment process. It accelerates drafting, research and assumption-testing, but introduces new risks around confidentiality, shadow usage, information reliability/quality and junior training. The practical lessons we hear from early adopters are to start with one relevant large task that fits the spec for an “AI-automatable task”, choose the best model to fit the task, build light governance around that task to test quality of AI-based execution, build a process of internal knowledge-sharing around that AI usage, and avoid cutting junior capacity too quickly, since juniors often drive adoption.

### What are OCIOs and investment consultants currently offering?

Publicly announced AI adoption by firms such as Mercer, Aon, Callan, WTW, Russell Investments, NEPC and Cambridge Associates falls into three buckets:

- **Internal GenAI LLM adoption**: Firms are using existing GenAI LLMs (e.g., ChatGPT, Claude, Grok, etc.) to improve manager, asset class and macro research workflows, document handling, knowledge retrieval and general productivity on all workflows.

- **Analytics platforms**: These support scenario modelling and portfolio construction decisions and are usually presented to consulting clients as advanced analytics rather than AI, even where machine learning is used at specific layers within the workflow (e.g., data extraction, providing a chat interface, etc.).

- **Manager due diligence**: AI and data science are being used to scale the number of managers that can be covered and improve insights, but these capabilities are rarely offered to clients as standalone tools and only used internally.

The implication is straightforward. Consultants are using AI to improve the advice they give, not to give clients the tools to do the work themselves. For allocators that want to make the best use of emerging AI tools, there are three generic options: 1) build internally using GenAI LLMs, 2) adopt specialist third-party AI tools, or 3) deepen reliance on legacy end-to-end platforms which are effectively integrating AI into their platforms. The best solution varies by phase of AI adoption and by use case as we describe below.

## How allocators should think about building their own AI capabilities

Allocators face two distinct architectural decisions, on different timescales. The near-term decision is which AI tools to adopt and which workflows to build with GenAI LLMs today (Phase 1). The longer-term decision is where and how to integrate a multi-agent AI architecture over the next two to three years (Phase 2). Phase 2 is a more holistic organisational overhaul, embedding agentic workflows into the firm's architecture at every level: from saving analysts time crunching manager data rooms, to integrating AI agents as IC members.

Phase 1 decisions are ideally made on the path to Phase 2. Every tool adopted, workflow built and API integrated in Phase 1 is a candidate to become a component the Phase 2 agent layer will later access. Canoe's data feed becomes what a capital-calls-agent reads from. Claude-based macro scenario workflows become what a portfolio construction agent invokes. Aladdin APIs become what a portfolio monitoring agent queries. Phase 1 is the focus of this paper. Phase 2 is the destination on the same roadmap.

## Phase 1: Three approaches to integrating AI today

Allocators have three approaches to Phase 1. Building with GenAI LLMs delivers judgement-led analytic capability at near-zero incremental cost. Consolidating on a Legacy End-to-End Platform provides an integrated operating system with AI features layered on top by the vendor. Buying Specialist AI Tools solves specific high-effort problems neither of the others handles well. These approaches are not mutually exclusive. Most allocators will run all three in parallel, deploying each where it fits the workflow best.

### Approach 1: Build with GenAI LLMs

GenAI LLMs such as ChatGPT, Claude and Grok are already highly capable tools for institutional asset allocators. For a range of use cases, a disciplined team can replicate a meaningful subset of the functionality offered by specialist AI tools, often at relatively low incremental cost.

Modern LLM platforms include features that allow allocators to build lightweight, repeatable workflows on top of the base model. For example, Claude Projects and ChatGPT Custom GPTs enable teams to preload reference materials—such as investment policy statements, prior investment committee memos, and internal style guides—and run repeated queries against those templates which shape the output.

With support from in-house software engineers, organisations can extend these models into more customised internal tools. This can include codifying standard analytical workflows—such as performance attribution, macro scenario development, cash flow reconciliation, or investment memo drafting—into reusable scripts that generate structured outputs in formats such as Word, Excel, or PowerPoint.

In enterprise configurations, including API-based deployments and controlled environments, it is possible to significantly reduce the risk of confidential data leakage, provided appropriate security, access controls, and governance frameworks are in place. However, this requires deliberate implementation and cannot be assumed by default.

Additionally, integration layers (e.g., APIs or model context protocols such as MCP servers) can connect LLMs to internal data sources and systems, such as document repositories, email platforms, or CRM systems, enabling more context-aware and operationally useful outputs. An **MCP server** refers to a server implementing the **Model Context Protocol (MCP)**—a relatively new standard for connecting AI models (like ChatGPT or Claude) to external tools, data sources (e.g., Bloomberg, Preqin, Factset, Capital IQ and Pitchbook), and systems in a structured, secure way.

Taken together, these capabilities allow a small, technically capable team to build a bespoke analytical layer tailored to specific workflows. For narrowly defined use cases, this can approximate some of the functionality of third-party specialist AI tools, often with greater flexibility and at lower cost. However, specialist platforms may still offer advantages in areas such as data integration, model validation, regulatory compliance, and large dataset ingestion.

Obviously, the greater the customisation, the greater the software development effort and the greater the utility. The choice ranges from leaving the base model untouched and using it as an off-the-shelf solution, through fine-tuning the model on proprietary data, to building a deeply customised proprietary model. As general guidelines, **Exhibit 1** sets out the three levels, the typical LLM engines, the degree of customisation involved, and the organisation size and investment level at which each level becomes appropriate.

**Exhibit 1: Three levels of GenAI customisation; most institutional allocators seek to land on Level 2**

| Level | What changes | What the organisation does | Example | Typical technologies | First-year cost | Common adopters |
|---|---|---|---|---|---|---|
| **1. Off-the-shelf GenAI LLMs** | Nothing — public model used as-is | Uses a consumer or enterprise LLM out of the box. Prompting and custom GPTs / Projects only; no integration with internal systems. | An analyst pastes a DDQ into Claude and asks for a structured summary. | ChatGPT, Claude, Gemini, Grok, Microsoft Copilot | <$250k | Most allocators today, including family offices and small endowments |
| **2. LLM rigged to your own data** | The inputs the model sees | Connects an off-the-shelf LLM to internal documents, CRM, email and accounting systems via APIs and Retrieval-Augmented Generation (RAG). The model itself is unchanged. | Claude connected to the firm's IC memo archive and CRM, so an analyst can ask what was concluded on a given manager in prior diligence. | Enterprise LLMs with API and RAG integration into internal data lake | $250k–$15M | Large endowments, foundations, sub-$50B pensions, regulated finance |
| **3. Proprietary / deeply customised models** | The model itself, end-to-end | Builds or heavily retrains proprietary models, hosted privately with custom architectures and a dedicated AI engineering team. | A SWF screening millions of filings per hour against proprietary signals, on infrastructure no data ever leaves. | Privately hosted foundation models on enterprise cloud infrastructure | $15M–$600M+ | Largest SWFs and pensions |

For a large university endowment, connecting an off-the-shelf LLM to internal data (Level 2) is the right level of customisation. Level 1 leaves the model disconnected from the firm's own documents: identical queries run across Claude, ChatGPT and Grok produce essentially the same output, because all three draw from the same public corpus. Without proprietary data, there is no source of differentiated insight. Level 3 is the territory of the largest SWFs and global banks, with $15M–$600M+ first-year costs and a dedicated AI engineering team, justified only where data sovereignty, query volume or latency requirements rule out a third-party API, which is rarely the case at endowment scale. Level 2, at $250k–$15M, gives an off-the-shelf frontier model access to the firm's own IC memos, manager database, CRM and accounting data via RAG and MCP, moving from generic productivity tool to a system that can answer institution-specific questions. This is the level the largest endowments, foundations and sub-$50B pensions are converging on.

There are also specific use cases where GenAI LLMs are not the right tool, regardless of governance. They can process a handful of documents on request, but not hundreds of GP reports on a recurring quarterly schedule. They cannot deliver the 95%+ extraction accuracy and full audit trails that accounting and audit require. They cannot conduct advanced factor analysis of a portfolio. This is where Approaches 2 and 3 may be necessary.

### Approach 2: Consolidate on a Legacy End-to-End Platform

Multi-module operating systems (**BlackRock Aladdin, SimCorp, Allvue, Dynamo**) run the full investment workflow on one stack: research, portfolio construction, risk, compliance, accounting and reporting. The trade-off is breadth versus depth. Allocators get consistency, a single source of truth and integrated AI features as the vendor rolls them out, but accept that any individual capability will likely lag a Specialist AI Tool, and that the AI roadmap is set by the platform, not the firm.

### Approach 3: Buy Specialist AI Tools

Single-purpose, AI-native products designed to solve a specific problem that neither a GenAI LLM nor a Legacy End-to-End Platform handles well. The best of these are narrower in scope but deeper in automation and scale within their targeted workflow area — **Vantager, Hebbia** or **Clade** for manager due diligence at full data room scale, Canoe Intelligence for post-investment data infrastructure.

### Four high-impact Phase 1 use cases where AI can add value

We focus on four areas where AI could add value for institutional allocators. The best tool for each use case is mapped out below in **Exhibit 2** crossing GenAI LLMs, legacy internal platforms and third-party specialist AI tools.

- **Manager due diligence and data room mining:** Ingesting GP data rooms and related documents, extracting and reconciling key information across sources, and producing structured outputs to support manager evaluation.

- **Portfolio construction and scenario-based asset allocation:** Defining macro scenarios, mapping them to portfolio exposures, and translating them into allocation decisions and ongoing monitoring. This can be used for long-term strategic asset allocation or shorter-term tactical allocation moves.

- **Workflow Management:** Managing the end-to-end investment process, including operation and investment team tasks including research, portfolio construction, trading, risk, compliance, accounting and reporting within a unified operating system.

- **Data infrastructure:** Collecting, extracting, standardising and distributing data from GP reports and other sources into downstream analytics, CRM and accounting systems.

**Exhibit 2: AI tools landscape for institutional allocators**

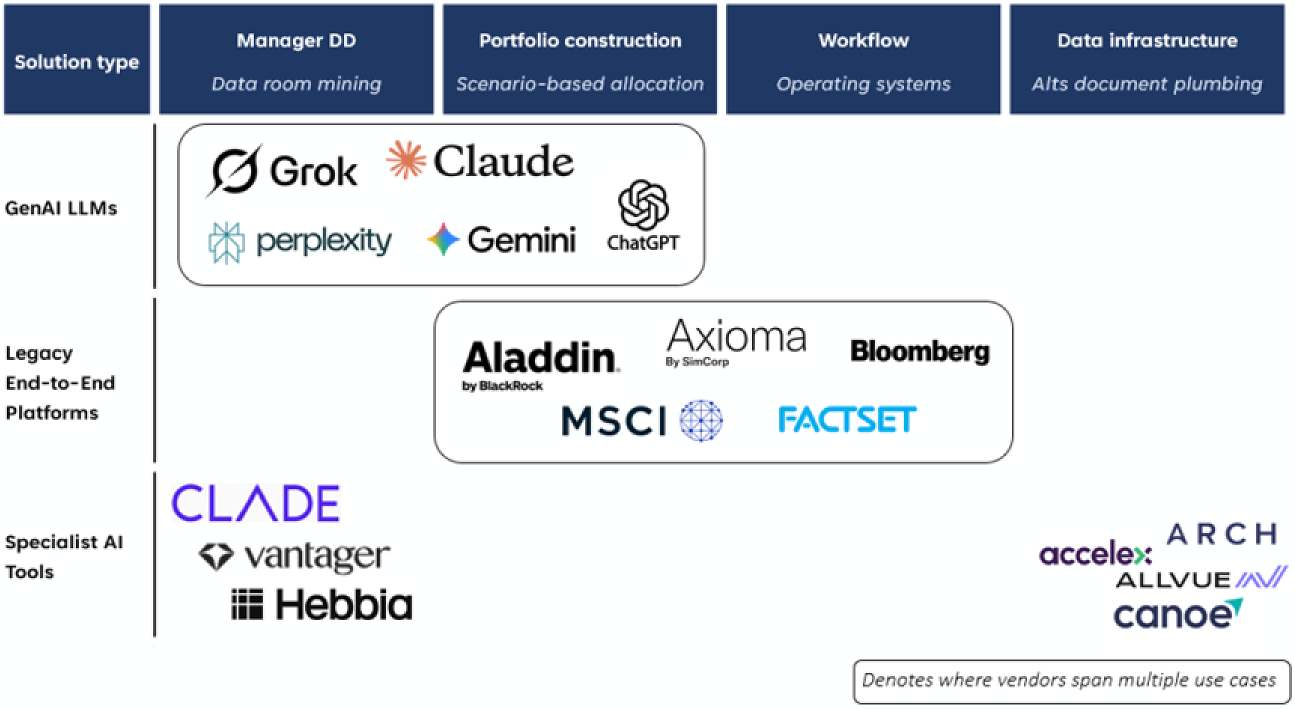


**Source:** Company disclosures; Vendor demos (April 2026)

## Phase 2: Integrating AI agents into firm architecture

**An end state that CIOs should seek is a multi-agent architecture in which AI agents serve as intermediaries between users and the underlying AI stack—the integrated set of models, data systems, workflows, and infrastructure that powers the firm's investment process.** The clearest articulation to date is MacArthur Foundation's "CIO Operating System" (Olusanya, April 2026), which proposes 15 purpose-built AI agents acting as virtual team members across the investment office. Drawing on the MacArthur model, **Exhibit 3** sets out TNI's view of such an end-state architecture as four layers: a data lake at the base, an analytics layer of point solutions above it, an agent layer of five to ten purpose-built agents, and CIO workflows at the top.

In the institutional investment context, an **AI agent** is a continuously operating software system that uses AI models, data access, memory, and decision logic to monitor conditions, interpret new information, and take or recommend actions in pursuit of defined objectives. What distinguishes an agent from a traditional chatbot or one-time AI query is not necessarily full autonomy, but **persistent operation**: it is typically "always on," continuously ingesting market data, news, portfolio information, and internal research, then updating assessments or triggering alerts as conditions change. An agent may be autonomous to varying degrees, but the defining characteristic is usually that it operates as an ongoing process with memory, monitoring, and workflow capabilities — rather than a tool that is manually invoked once for a single prompt or analysis.

In Step 1 of the agentic model development, users and applications access the data and analytics layers directly, with each tool supporting workflows. In effect, users manually bridge from Layer 3 (applications) to Layer 1 (CIO workflows), while agents undergo development.

In Step 2, the agent layer intermediates: agents compose responses across the data lake and applications, enabling workflows to be conducted agentically. The most powerful, alpha generating AI Agents for institutional investors must be those which actively monitor the risk of the portfolio –

at the overall portfolio level, asset class level, manager level and, ultimately, at the individual asset level. In Exhibit 3, we show 10 of what could be 15 or more agents but have chosen those we believe could be most value added. These include a Portfolio Risk Monitoring Agent, Manager Risk Monitoring Agent, Live Macro Scenario Agent (with research aggregation), Manager Sourcing Agent and a IC Bull / Bear Member Agent. Other more operationally oriented examples include a Portfolio Report Generation Agent, Manager Onboarding Agent, Compliance Monitoring Agent, Capital Call Management Agent and a Liquidity Monitoring Agent that pressure-tests live investment theses.

Before building toward this, allocators should be realistic about the foundation it requires. The agent layer is only as good as the data lake beneath it, and most institutions remain behind in turning the data layers into the “sole source of truth” enriched with universes of unique data. Even advanced practitioners report data hygiene as the binding constraint on getting value from their AI tooling and infrastructure.

**Exhibit 3: Example Phase 2 agentic architecture for an institutional investment office**

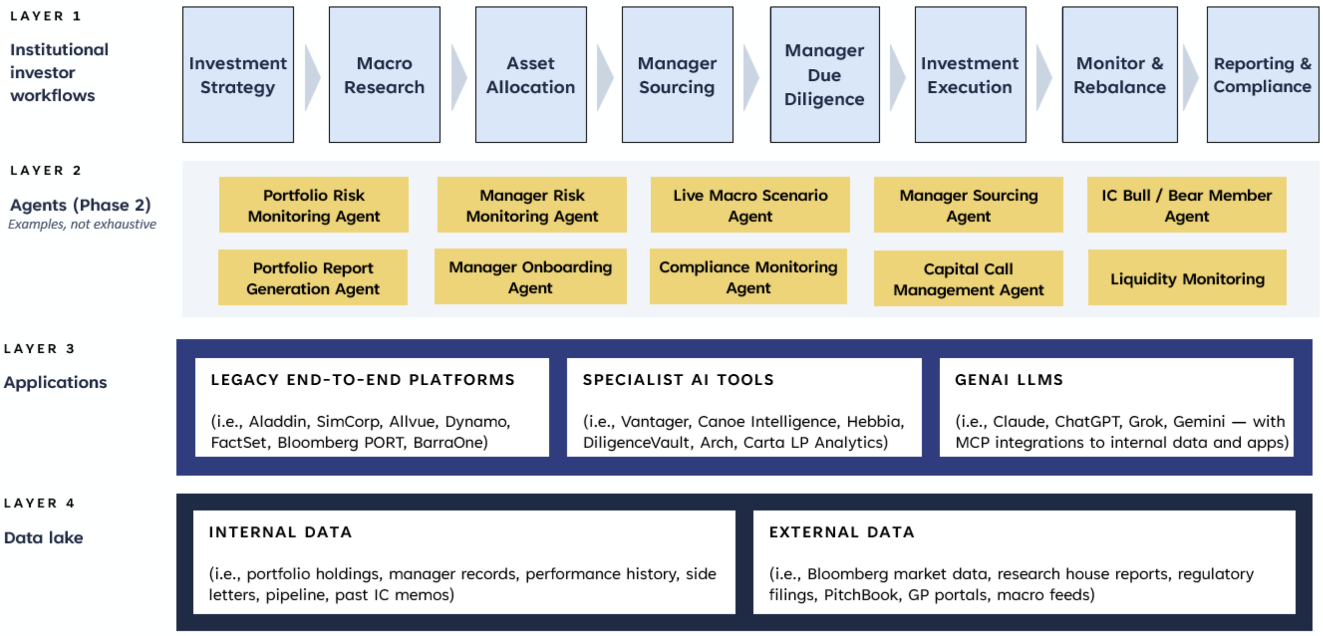


The Phase 1 decisions analysed in the rest of this paper are the building blocks of this architecture. Adopting **Canoe** builds the post-investment data layer. Adopting **Vantager, Hebbia** or **Clade** builds the manager research feed into the applications layer. Integrating **Aladdin** via APIs better informs portfolio construction and risk management. Connecting an off-the-shelf LLM to internal data via RAG and MCP is the precondition for any agent layer that follows. Each Phase 1 tool delivers standalone value today and contributes a specific component to the Phase 2 architecture. Allocators that move now will have the data, applications and integration foundations substantially in place within two to three years. Those that delay are just likely to start the same build later.

Peer AI-forward GP firms are emerging as a useful external learning channel during this build. Several institutional allocators now hold quarterly calls with AI-native managers who share reading lists, comment on which frontier models are outperforming where, and explain how they are reallocating usage across Claude, ChatGPT and other models as the frontier shifts. For LPs without dedicated AI engineering teams, partnering with AI-forward GPs offers a low-cost way to track tool selection, prompt design and architectural choices.

## Phase 1 Use Cases

### Use Case #1: Manager due diligence and data room mining

This is one of the few areas where AI already delivers clear, practical value that can translate into significant alpha improvement for asset allocators. Third party specialist AI tools can materially reduce diligence time, improve consistency, and surface issues that are easy to miss in a manual process. GenAI LLMs cover a meaningful share of the work for smaller teams but break down when due diligence extends to increasingly enlarged GP data rooms and where diligence runs across multiple managers for a given mandate. Such tools may be essential to an institution's migration from the endowment model to TPA (Total Portfolio Approach) to the extent TPA has us comparing managers across asset classes on the basis of overall contribution to portfolio alpha and alpha volatility across the total portfolio."

**The ideal automated manager due diligence system would ingest a full GP data room**, including pitch deck, DDQ, audited financials, LPA, side letters, quarterly letters and track record. The average private equity VDR contains between 500 and 2,000 discrete documents. Existing automated manager DD tools combine that data with LP-generated analysis and convert the lot into a structured, queryable dataset, and ultimately into the investment memo (IM) to the investment committee. The tool provides the first pass IM and augments iterations over the DD process, but what ends up in the investment committee's email boxes is the combination of AI and human, but clearly 100% owned by the human presenting it. AI models like **Vantager, Hebbia** and **Clade** reconcile inconsistencies across documents and surface a ranked work queue showing what has changed, what conflicts, what is missing, and what requires follow-up. It would also monitor external signals such as Form ADV updates, litigation, press coverage and team movements, maintaining a single, current view of each manager.

Investment memo output matters as much as DD data ingestion. The system should be able to draft the majority of an IC memo along pre-defined lines, producing a customised, consistent, decision-ready document covering the fund's investment strategy, investment process, track-record, team and terms. In practice, this means generating a strong first-pass memo, allowing teams to focus time on judgement rather than assembly, and to stop early on managers that should not progress based on this "shallow deep dive", drawing initially on data room inputs.

**What can be done today with Claude or ChatGPT is already substantial.** A team with a well-organised document set can extract fund terms, summarise track records, flag obvious inconsistencies between decks and DDQs, draft first-pass IC memos and generate follow-up questions. The same workflow can be turned inward: connecting an LLM to the firm's own archive of investment memos, post-mortems and IC minutes can surface recurring flags that preceded poor manager outcomes and patterns in historical alpha, building pre-mortems on prospective managers from the firm's own past mistakes. For an allocator reviewing any number of new managers a year, this covers a large share of the workload, increasing the number of managers that can be reviewed at a higher level of initial input than our historical "first passes" allowed.

Pulling data out of documents has been somewhat commoditised by GenAI LLMs. The harder task is sorting data into organised and consistent formats <u>across data sources</u> and checking that data against itself. This is where Specialist AI Tools provide a real advantage.

TNI has integrated GenAI LLMs into manager sourcing, screening and due diligence to increase speed and deepen the quality of analysis. In practice, we use AI to support first-pass outputs including performance breakdowns, beta-adjusted alpha, sector and geographic attribution, return concentration, position-sizing skill, and basic statistical tests of significance and persistence. This allows us to move quickly from raw materials to a structured view on where returns are coming from and whether they appear repeatable. The process remains analyst-led: an Associate collects all of the relevant data from the manager and external sources and engages the LLM to translate these into our standard DD report in sections: strategy, team, process, performance and terms. Outputs are not taken at face value. This is where the real work begins. The LLM output is used to frame and elevate the initial internal discussion, focusing time on the key judgement calls rather than on assembling the underlying analysis. This does require you to clearly instruct the LLM on the desired output format.

The output we now receive from AI LLM supported due diligence demonstrates that a disciplined GenAI LLM workflow can produce a much higher quality first-pass review than would have been practical in the past, including the manager's strategy, performance, key risks, and key follow-up questions for fund managers. This materially accelerates the first stage of diligence by moving quickly from raw materials to a structured view that an investment professional can challenge, refine and take forward or to discard.

The limitations become clearer in larger, more institutional settings. GenAI LLM workflows can support individual pieces of analysis, but they are harder to run consistently across large data rooms, multiple managers and "DD refreshes" years after the initial investment. The standard LLM models do not automatically create a permanent record of prior reviews, maintain a clean audit trail, benchmark new opportunities against a structured history, or monitor changes in manager disclosures and regulatory filings over time. To overcome these shortcomings, we have two choices. One is to have internal or external specialist software developers to connect off-the-shelf LLMs to internal documents, CRM, email and accounting systems via APIs and Retrieval-Augmented Generation (RAG). The other option is to procure existing tools that solve these problems and integrate your data and external data like Pitchbook, Preqin and Capital IQ to a manager due diligence tool. Nothing reliable existed to perform in this way just two years ago. Today, we recommend institutions have a look at **MSCI's Vantager**, **Hebbia** and **Clade** which stand out ahead of others, based on our review to date.

**Vantager**, acquired by MSCI in March 2026 and now part of the **MSCI Diligence Platform**, is built for institutional-scale manager diligence. It was under construction for two years prior to launch in June 2025 and today is used by just 15-20 allocators representing $100B of AUM. The platform ingests full GP data rooms, typically 900–1,200 documents, by zip upload or email pipe, then classifies each page and maps content into structured investment, operational and legal diligence fields and reports. It can generate IC-ready memos using the allocator's own templates. Its main advantage is the quantum of documents it can ingest. It also offers a cross-manager database growing over time within the allocator's proprietary version of the tool. Every uploaded data room remains queryable, including passed deals, enabling side-by-side comparison across managers and vintages. MSCI ownership is expected to add stronger private-asset data, index data and benchmarking. It can be used on managers across all asset classes. Its key strength is standardisation and reconciliation across large numbers of manager data rooms. No current client references have been taken yet, and we have not tested Vantager in our own data rooms. **See Appendix for more detail on the tool.**

**Hebbia** is the most credible alternative to Vantager for the data room-mining problem but solves it differently. Where Vantager is purpose-built for LP manager diligence with structured templates and reconciliation logic, Hebbia is a general-purpose AI research platform: a user uploads hundreds or thousands of documents — data rooms, filings, transcripts, manager letters — and runs structured queries that return source-linked answers across the whole corpus in parallel. Asking "summarise fund terms, fees and key-person provisions across these 30 LPAs" returns a cited table in minutes. It claims adoption across more than 40% of the largest asset managers by AUM, with named users including KKR, Centerview Partners and MetLife. For the manager DD use case specifically, the limitation is that Hebbia's customers have been managers and banks, not allocators. There are no LP-specific templates, DDQ workflows, manager database or regulatory filings monitoring of the kind Vantager provides out of the box. They do provide API access to relevant external data bases like Preqin, Pitchbook, Snowflake and FactSet. The company is now expanding into the institutional investor segment and building allocator-specific capabilities, but this is early and there is no public LP reference base yet. **See Appendix for more detail on the tool.**

**Clade** is another alternative to Vantager; an AI-native platform built for the LP side of the manager diligence workflow. Where Vantager originated in manager diligence and Hebbia in general-purpose research, Clade was built specifically for allocators. It ingests an allocator's document set — decks, fact sheets, PPMs, LPAs, side letters, emails and attachments, both historical and ongoing — and organises it into a queryable, source-cited set of per-fund workspaces. From that corpus it runs side-by-side fund comparisons, extracts structured data points from LPAs, drafts IC memos, tear sheets and ODD reports in the allocator's own format, and answers natural-language queries with citations to source. It routes queries across several LLMs to manage cost. Its main limitation is that Clade's edge rests on allocator-specific workflow design rather than proprietary data or models and Clade lacks the proprietary private-markets datasets and benchmarking that Vantager is planning to integrate from MSCI. Named clients include Summation Capital, Purdue Research Foundation, SCERS and the University of Florida. We have not tested Clade in our own data rooms or taken client references. **See Appendix for more detail on the tool.**

Our comparison between the MSCI Diligence Platform, Hebbia and Clade is focused very much on their future development potential. That is a bet between MSCI as a large (bureaucratic?) public company with huge resources and databases, the investors in Hebbia being Peter Thiel and Andreessen Horowitz, and Clade as the most allocator-focused of the three but also the smallest and least well-resourced. Who will develop the fastest is not clear.

Other manager DD vendor names that we reviewed include **Finpilot, Rogo, Opto Investments and DiligenceVault** which have broader scope across the investment process, with partial manager DD functionality relative to Vantager, Hebbia and Clade. None of these offer full data room ingestion and reconciliation. Several are GP solutions, not LP. The closest appears to be **DiligenceVault,** which is not AI-enabled but serves the same purpose as Vantager of making data room extraction easier, by getting managers to upload their data rooms, and extracting data in a structured format to a centralised database which LPs can subscribe to.

**Manager Selection Algorithms.** Over time, this universe of manager data should contain hundreds of manager due diligence exercises, creating a structured and continuously expanding dataset from which statistically meaningful patterns can be identified. Agentic AI systems can apply machine learning, pattern recognition, natural language processing, and probabilistic modelling techniques to analyse relationships between manager characteristics, portfolio behaviour, market regimes, and subsequent investment outcomes. By continuously monitoring and learning across both structured data (performance, exposures, factor sensitivities, turnover, liquidity) and unstructured data (manager letters, meeting notes, DDQs, transcripts), AI can identify recurring attributes associated with persistent alpha generation. Unlike traditional one-time analytical tools, these systems can operate continuously, updating their assessments dynamically as new information and macro conditions emerge. Their advantage is not that they "predict" markets independently, but that they can ingest, compare, and evaluate vastly larger volumes of information than human teams alone, while applying a more systematic and less behaviourally biased analytical framework.

**AI tools in manager selection may become a competitive advantage for the large private markets investment platforms (fund of funds) challenging the value added of internal manager selection teams.** An important observation we have on AI used in private asset manager selection is that the largest private-markets allocators, including fund-of-funds in particular, increasingly are building up a structural data advantage that institutions below, say, $50B in AUM, will struggle to replicate. Firms such as HarbourVest, Hamilton Lane, StepStone and AlpInvest can access granular portfolio-company revenue, earnings and balance sheet data across thousands of private holdings during any given GP's ownership period. This allows these scale private markets investors to more accurately measure whether a manager's historical investment returns reflect genuine company selection and post-acquisition value-add, or simply exposure to sectors that performed well (e.g., luck). GenAI LLMs do not create this advantage, but they make the data easier to structure, query and compare at scale. Whoever has the largest fund and deal database has inherently better ability to assess real manager value add vs luck. This advantage has theoretically sat with whichever investor had the broadest base of funds and underlying portfolio company data in the past. Only today, with the power of AI, can such investors fully exploit that advantage. This highlights the importance of any private markets investor seeking to mine their own company level databases in manager due diligence. Small to medium sized institutional investors will need to think more carefully about the internal vs outsourced private markets fund research question in the name of alpha generation.

**A note on confidentiality and data security.** Almost all Specialist AI Tools operate within enterprise-grade security perimeters and contractually exclude client (asset manager data room) data from model training. The more important question is architectural: whether the vendor processes documents on its own infrastructure or routes them through a third-party LLM API. Some specialist AI tools build their own models trained on alternative-investment documents; others sit on top of OpenAI or Anthropic APIs, which raises additional vendor-management and data-residency questions. CIOs should require each vendor to disclose this explicitly, alongside SOC 2 Type II certification, encryption standards and access-control logs, before sharing confidential GP materials.

### Use Case #2: Portfolio construction and scenario-based asset allocation

This is the least mature of the four use cases. No current AI offering automates or replaces human judgement on the full workflow supporting overall portfolio construction (macro research, scenario generation, strategic asset allocation, manager position sizing, risk management and tactical asset allocation). The most alpha-driving activity under the heading of portfolio construction is strategic asset allocation, where best practices start with scenario definition and work through expected returns, risk and correlations to arrive at the optimal long-term asset allocation. GenAI LLMs are today remarkably useful for automating what has been a labour-intensive analytical effort, but one that is usually only undertaken once a year. What AI allows, that has not been practical in the past, is to revisit the optimal asset allocation on a more frequent basis, as macro views change with geopolitical, political and purely economic events -- i.e., AI enables investors to deploy a more dynamic or tactical asset allocation process during the year, that may translate into meaningful alpha.

**The ideal system for deploying a more dynamic asset allocation strategy** would start from the total portfolio's full look-through exposures across all assets, looking through to the underlying holdings of all asset managers. Once AI is working in the data layer of the business that monitors manager holdings, look-through exposure analysis is possible to do on a more frequent basis, and in response to macro events. The AI-enhanced model would prompt the CIO to define a small number of macro scenarios from a given point in time and forecasting to a twelve to thirty-six-month horizon. It would then assign explicit probabilities to each of those scenarios or update those probabilities. It would translate each scenario into asset-class return/risk/correlation outcomes, recommend a scenario-weighted portfolio with specific actions to add, trim, hedge or fund, and track drift against that plan. The output would be a decision-ready IC memo, explaining the recommended change in asset allocation in both qualitative narrative and quantified factor terms.

In practice, no platform does this today. Legacy end-to-end platforms (e.g., MSCI BarraOne, Bloomberg PORT, etc.) handle the data availability side well but leave the core judgement step untouched. Deciding which scenarios matter, assigning probabilities, and updating them over time remains a highly manual analytic process, overlayed with large doses of judgement. This is where the LLMs come in quite usefully and adequately. Value add is down to unique insights into macro factors and asset class dynamics.

GenAI LLMs are highly capable here. Prompting a model to generate a small set of macro scenarios, assign rough probabilities, and outline an allocation for each produces output that is credible as a starting point for IC discussion. **Exhibit 4** shows Claude output from a simple query for three macro scenarios with expected returns and recommended asset allocation.

**Exhibit 4: Claude output from simple query for downside, base case and upside scenarios**

| Section | Metric / Asset Class | Downside | Base Case | Upside |
|---|---|---|---|---|
| **Macro assumptions** | Probability | 20% | 60% | 20% |
| | Scenario label | Low long-term return world | Normalised return world | High long-term return world |
| | Global GDP growth | 1.5-2.0% | 2.5-3.0% | 3.0-3.5% |
| | Inflation | 2.5-3.5% | 2.3-2.8% | 2.0-2.5% |
| | US 10Y Treasury yield | 3.0-3.75% | 3.75-4.25% | 4.25-5.0% |
| | IG credit spreads | ~90 bps | ~120 bps | ~180 bps |
| | HY credit spreads | ~300 bps | ~400 bps | ~600 bps |
| | HY default rates | 2-3% | 3-4% | 5-6% |
| | Corporate margins | Elevated / peak | Stable | Depressed, improving |
| | Equity valuations | Expensive | Fair | Cheap |
| | Private market valuations | Expensive | Fair | Attractive |
| | Globalisation / geopolitics | Managed stability; policy support | Regional fragmentation | Post-reset stabilisation |
| | | | | |
| **Asset class return assumptions** | Developed equities | 4.00% | 6.50% | 9.50% |
| (avg next 5-7 years) | Emerging equities | 5.00% | 7.50% | 11.00% |
| | IG credit | 3.50% | 4.50% | 6.00% |
| | HY credit | 4.50% | 6.50% | 8.50% |
| | Government bonds | 2.50% | 3.50% | 4.50% |
| | Absolute return hedge funds | 4.50% | 5.50% | 6.50% |
| | Private equity | 7.00% | 10.00% | 13.50% |
| | Venture capital | 6.00% | 11.00% | 16.00% |
| | Private credit | 5.50% | 7.50% | 9.50% |
| | Real assets | 5.00% | 6.00% | 7.50% |
| | Cash | 2.00% | 2.50% | 3.00% |
| | | | | |
| **Optimal asset allocation** | Developed equities | 18% | 27% | 28% |
| | Emerging equities | 4% | 7% | 10% |
| | IG credit | 22% | 11% | 5% |
| | HY credit | 5% | 5% | 7% |
| | Government bonds | 18% | 8% | 5% |
| | Absolute return hedge funds | 5% | 6% | 4% |
| | Private equity | 13% | 15% | 18% |
| | Venture capital | 3% | 5% | 8% |
| | Private credit | 5% | 7% | 8% |
| | Real assets | 5% | 7% | 5% |
| | Cash | 2% | 2% | 2% |
| | Total illiquid allocation | 26% | 34% | 39% |
| | Expected portfolio return | 4.30% | 6.80% | 10.00% |
| | Expected portfolio vol | 7.5-8.5% | 10-11% | 12-13% |

As we evolved this use case, we had Claude interview us about our qualitative views on each asset class. It then provided us with a recommended optimal asset allocation in line with the probabilities of each scenario. We also input our own asset class return assumptions as a basis for the model varying those in different scenarios. This exercise was run on asset classes, but it can easily be run on betas and factors. Obviously, all assumptions have to be challenged. So the "AI value add" is in not having to build and maintain one's own model. To the extent we have trained the LLM with our learning about macro-scenario relationships to asset class returns, asset class interactions, liquidity constraints, etc, we have developed an AI LLM integrated macro model. For example, the scenario updates will draw on the external and internal universe of expert data on macro outcomes (e.g., inflation, rates, etc.), look at current asset class prices and asset class challenges and factor those in without the internal team having to update each input.

The AI vendor landscape for overall portfolio construction is dominated by legacy end-to-end platforms rather than specialist AI Tools. **BlackRock Aladdin** remains the clearest example of an end-to-end platform supporting portfolio construction. It brings together **portfolio construction, risk management, trading, and operations** into one integrated platform. It is used by more than 240 institutions representing over $22T of AUM, and its AI layer now includes **Aladdin Copilot**, **eFront Copilot** for private markets and the **Thematic Robot** for basket construction. Built on Azure OpenAI with an agentic plugin architecture across roughly 100 front-end applications, these tools provide natural-language access to Aladdin's underlying APIs, automate reporting and reduce team-member onboarding friction. The important constraint is that Copilot operates within verified Aladdin data and is explicitly non-advisory. It improves how users access, query and navigate the platform, but it does not independently generate or recommend asset-allocation changes. **Exhibit 5** breaks down the portfolio construction process into it six steps and indicates where the legacy end to end platform is best augmented with inputs from GenAI LLMs. The key value added from LLMs is macro scenario generation. This flags a need to work with your macro research providers (independent research houses like BCA and macro hedge funds like Bridgewater), as they may well be duplicating some of your efforts here and can in any case become inputs to your models. Firms like Two Sigma, DE Shaw and AQR are undoubtedly further down this path than we are.

**Exhibit 5: Portfolio Construction and Asset Allocation: GenAI LLMs and Legacy End-to-End Platforms are complements, not substitutes**

| Requirement | Description | GenAI LLMs | Legacy End-to-End Platforms |
|---|---|---|---|
| Total-portfolio view with look-through exposures | Maintains a governed, current view of the total portfolio across public and private assets, with look-through to underlying holdings and factors. | No native portfolio data layer; reasons only on what is in the prompt. | Core capability across all platforms. |
| Macro scenario generation and framing | Generates plausible 12–36 month scenarios with transparent, challengeable reasoning. | Produces credible scenario sets quickly, suitable as a starting point for IC discussion. | Not a generative function; scenarios must be defined manually. |
| Translation of scenarios into priced outcomes | Converts scenarios into expected asset-class returns, factor moves and risk impacts. | Asset-class-level impacts only; cannot reprice live holdings through a calibrated factor model | Core strength — factor models and stress engines price scenarios across asset classes. |
| Scenario-weighted portfolio recommendation | Recommends a probability-weighted portfolio with specific add, trim, hedge or fund actions. | Runs probability-weighted optimisation across asset-class mixes; cannot reflect live holdings or portfolio-specific constraints | Constructs optimised portfolios across live holdings and constraints; requires user-defined scenarios and probabilities as inputs |
| Drift monitoring and risk attribution | Tracks portfolio drift versus plan, decomposes risk, and flags re-balancing needs. | No portfolio memory between sessions; cannot track drift | Continuous monitoring, factor decomposition and risk attribution are core functions. |
| IC memo and decision support | Produces a cited, IC-ready output explaining the allocation in both narrative and factor terms. | Strong on narrative; weak on factor explanation and governed audit trail. | Strong on factor reporting; weaker on narrative output. Aladdin's Auto Commentary is an early step. |

Key: Strong Average Weak

**Source:** Vendor websites; Vendor meetings; Vendor demos

Beyond BlackRock Aladdin, other end-to-end portfolio management platforms include **MSCI BarraOne**, **SimCorp Axioma, FactSet** and **Bloomberg PORT**. These remain powerful quantitative platforms for factor modelling, risk attribution, stress testing and portfolio analysis. Natural-language interfaces may make them easier to use, but the core value still sits in the existing factor models, risk systems, stress-testing tools and portfolio analytics. Because the AI layer is being built on top of these platforms, the pace and quality of integration will increasingly differentiate vendors that have looked broadly similar for years. CIOs should treat the vendor's credible AI roadmap as a first-order selection criterion, monitoring which platforms are shipping useful AI functionality fastest, where it is genuinely embedded in the investment workflow versus bolted on for marketing, and how quickly each vendor closes the gap to the agentic, natural-language standard now being set by Aladdin.

## Use Case #3: Workflow Management

Workflow platforms bring portfolio construction, risk analysis, trading, compliance, accounting and performance reporting into a single operating system. This is the operational backbone after the investment decision is made. **BlackRock Aladdin i**s the archetype, with **SimCorp** and **FactSet** as credible alternatives. These systems are mature, powerful and expensive, and have been embedded in institutional workflows for decades, usually customised by the user. AI is turning platforms like Aladdin, SimCorp, and FactSet from systems you operate into systems you interact with conversationally and automate intelligently.

The AI story for workflow management is recent and gradual. AI is being added to existing functionality rather than changing what these platforms fundamentally do. **Aladdin Copilot** provides natural-language query, faster onboarding and report automation. BlackRock **eFront Copilot** extends similar functionality to private markets. **Auto Commentary** generates client-ready narratives in Aladdin Wealth. The practical value is productivity improvement: faster access to underlying data, fewer manual report builds and lower training burden for new users. This is useful, but it is not new investment intelligence. Copilot is confined to verified Aladdin data and is explicitly non-advisory. It does not recommend allocation changes, identify novel signals or replace investment judgement.

Exhibit 6 breaks workflow management into its 6 source of value add and compares Aladdin, Simcorp One and FactSet against each other on those 6 dimensions. The differences are not stark, so incumbency wins.

**Exhibit 6: AI in Workflow Management – All seem to be doing what is available**

| Area | BlackRock Aladdin | SimCorp One | FactSet |
|---|---|---|---|
| **Core AI approach** | **GenAI "Copilot" layer over risk & portfolio data** | **AI embedded in operations & analytics workflows** | **AI embedded in data, research & user interface** |
| **Natural language interface** | Produces credible scenario sets quickly, suitable as a starting point for IC discussion. | Produces credible scenario sets quickly, suitable as a starting point for IC discussion. | Produces credible scenario sets quickly, suitable as a starting point for IC discussion. |
| **Workflow automation** | Assists navigation, analysis, and reporting tasks | Automates operational workflows (reconciliation, reporting, rebalancing) | Automates research, reporting, and data workflows |
| **Data ingestion (AI-driven)** | Limited direct focus (relies on structured data feeds) | Strong: extracts and processes operational data (docs, reports) | Very strong: ingests unstructured data (PDFs, filings, transcripts) |
| **Reporting & IC memo generation** | Drafts portfolio insights and explanations | Automates reporting and client outputs | Generates research summaries and reports |
| **Risk & scenario support** | AI-assisted interpretation of risk and scenarios | AI-enhanced optimisation and analytics | Limited vs others (more data-focused) |
| **Integration with internal systems** | Integrated within Aladdin ecosystem | Deep integration across front-to-back workflows | Expanding via APIs and MCP-style connectors |
| **Primary value of AI** | **Faster insight from complex risk data** | **Efficiency and automation of operations** | **Faster data access and research workflows** |

**Key:** Strong | Average | Weak

For the purposes of this paper, workflow is not a genuinely AI-led category. Most CIOs already use one of these platforms, and AI will continue to improve usability at the margin rather than displace the core systems. We include the category for completeness, but it is not where AI adoption decisions should be focused.

## Use Case #4: Data infrastructure

Data infrastructure is largely a solved problem, and one where dedicated non-AI native systems are required. GenAI LLMs are no substitute. The challenge is not how smart the model is, but whether the system can efficiently connect to the right sources, process data consistently and produce reliable outputs.

**The ideal AI-enabled data infrastructure system** must connect to GP portals, collect capital call notices and quarterly reports, extract key data such as NAV, IRR, cash flows and portfolio company metrics, put it into a consistent format, and feed it into the allocator's performance analytics, CRM and accounting systems. It must also reconcile manager-reported data with internal records, where differences often arise from timing, classifications, FX, valuation dates, fee treatment or cash-flow mapping. This matters because the same fund data can appear differently across GP reports, administrator feeds, accounting systems and portfolio analytics. AI can help classify documents, match entities, flag exceptions and extract data, but the model is not the important part. What matters is whether the system can capture data reliably across thousands of documents, reconcile it against internal records, and send clean outputs into the allocator's existing systems.

AI is now improving this process, but within a dedicated infrastructure system rather than as a standalone tool. In **Canoe Intelligence**'s case, AI helps classify documents, identify fund and entity names, extract data from inconsistent GP reporting formats, and flag items that require human review. This is especially useful in private markets because managers do not report in a standard format, and the same information is often buried in PDFs, capital call notices, quarterly letters and spreadsheets. The role of AI is to reduce the manual work needed to turn messy manager documents into clean, usable data. The important point is that AI works because it is embedded inside Canoe's ingestion network, data model, reconciliation process and downstream integrations.

GenAI LLMs are not competitive in this use case. They do not connect to GP portals, cannot run ongoing data collection processes, and cannot produce audit-grade outputs across thousands of documents.

**Canoe Intelligence** is the institutional standard. It connects directly to fund administrators and GP portals, ingests documents at scale, and extracts structured data into a central system. Coverage extends to 44,000+ private funds across 400+ clients, with over $8T of assets under automation and more than 50M documents processed annually. Canoe then sends clean data into downstream systems via APIs. The reported productivity gain, a 15-to-20-fold increase in funds processed per employee versus manual workflows, is material. For any allocator running a manual process across a meaningful private markets book, the return on investment is immediate. The April 2026 Canoe and Bloomberg partnership reinforces its position by embedding private markets data directly into Bloomberg PORT Enterprise, using FIGI as a common identifier, and making private fund data available alongside public-market analytics on the Terminal. More detail is provided in the Appendix.

The closest direct competitors are **Arch** and **Carta LP Portfolio Analytics, the rebranded Accelex platform** acquired by Carta in October 2025. Arch is built for a similar job, including K-1 collection, portal access and cash-flow tracking, but is positioned more towards family offices and wealth managers, where Canoe's pricing may be too high for the use case. Carta LP Portfolio Analytics offers a similar document-ingestion workflow, with a natural advantage for LPs whose GPs are already administered on Carta's network of 125,000+ LPs. **Allvue Systems** is a broader private-markets platform whose data extraction and reporting tools compete with Canoe where the buyer wants data collection, portfolio monitoring, fund administration and CRM in one system.

## Conclusion

Institutions are off and running with AI tools supporting their investing activities, but it is still very early days. We failed to find any institution or research suggesting anything like the Phase II description above, i.e., building a wall of integrated AI agents mapped onto every major investment process. But that could be just a few years away, given how fast the technology is developing. The key question today is what to outsource vs build in-house. We have provided a roadmap for that which outsources manager due diligence and data infrastructure tools to start, but counting on workflow and portfolio management platforms already used in-house, to be augmented with various AI capabilities in their offering to you.

**Manager diligence and manager monitoring is the clearest starting point:** GenAI LLMs can already support first-pass extraction, analysis and memo drafting for allocators reviewing a modest number of managers each year. Specialist AI Tools become more compelling as entire data rooms can be ingested. Such tools allow investment teams to screen a broader universe of investment choices and removes type II error where we screen out good managers in the interest of time, using more superficial screening mechanisms or rules like "no first time funds." We certainly subscribe to the investment principle of the more one puts in the funnel, the higher the quality of what comes out. Hence, more alpha.

**Portfolio construction requires a hybrid AI LLM / legacy platform model:** GenAI LLMs are useful for scenario generation, probability framing and IC narrative. Legacy End-to-End Platforms remain necessary for factor analysis, stress testing, risk attribution and implementation. Neither category yet delivers the full scenario-to-portfolio workflow. It gets interesting when these two can merge and proprietary data bases arm LLMs to provide timely insights on tactical asset allocation moves or major risk concentrations.

**Workflow Management is not a standalone AI adoption priority:** Legacy End-to-End Platforms will continue to add AI features, but these are mostly usability improvements. CIOs already using Aladdin, SimCorp or FactSet should benefit from those suppliers' AI-enhanced upgrades. We don't see anything out there that should replace legacy workflow management systems,... yet.

**Data-infrastructure AI Products offers huge efficiencies that off-the-shelf GenAI LLMs can't.... today:** Data infrastructure (aka "data layer" or "data lake") is the critical building block for all AI applications across the institution. Successful manager due diligence and monitoring and successful portfolio construction and risk management gain their greatest potential alpha enhancing capabilities through access to proprietary data being merged with external data. The human input to the data layer is the most critical, which is around identifying, accessing, and/or creating the most valuable inputs to investment insights. Canoe Intelligence and its peers solve a hard operational problem that GenAI LLMs cannot address on their own: connecting to GP portals, extracting data at scale, reconciling it against internal records and feeding clean outputs into existing systems.

The early, practical conclusion is to use GenAI LLMs where the task is still highly dependent on human judgement, but labour intensive in the initial data gathering/updating stages. AI's primary benefit is speed to our first-pass analysis. This applies most readily to manager selection, manager monitoring and overall portfolio construction and risk management.

**Culture is critical to AI integration.** The natural tendency for humans is to believe that we are always the better option to a robot, when the more objective among us know that the cyborg wins every time. The Terminator movies taught us this in the 80's. My message is that the culture matters hugely in driving AI integration and the philosophy ideally is one where the full team is always searching for the optimal mix and construction of human and machine. The larger obstacle is the portion of the team fighting AI usage, not the few that want to fully replace themselves with AI.

**No pure AI output should reach the investment committee.** Hallucinations and incorrect summaries remain a real failure mode of even the strongest models. Most CIOs allow AI to draft a first-pass memo, but require human augmentation, additional analysis and judgement before the document is decision-ready. Cross-LLM review, where a rival model checks the original's output, is a useful intermediate step but does not remove the need for human eyes.

**Caution: GenAI LLMs are rapidly improving.** Before procuring Specialist AI tools that generally have huge switching costs, explore the likely future developments with Anthropic, OpenAI, Google and X. A shortcoming of this paper is that we have not mapped out the likely future functionality of these rapidly improving tools.

We finish with **Exhibit 7** summarising how GenAI LLMs and Specialist third-party AI tools can be used today to generate more alpha from institutional portfolios. TNI has not directly tested each tool referenced. Views are based on vendor discussions and AI-assisted research.

**Exhibit 7: AI value-add by use case and tool type**

| Use case | GenAI LLMs | Specialist AI Tools |
|---|---|---|
| **Manager diligence** | Useful for allocators reviewing <50 managers a year. Strong on extraction, summarisation, first-pass IC memos and follow-up questions. Less effective at institutional scale | MSCI Vantager may well be alone in its ability to ingest entire data rooms, augment internal analysis and other DD inputs into one investment memo. Teams can look at far more managers to find the gems. |
| **Portfolio construction and scenario allocation** | Strong for scenario generation, probability framing and narrative construction. Not yet capable of decomposing live portfolios into factor exposures, running optimiser-backed allocations, or maintaining a governed record of decisions and outcomes. | No Specialist AI Tool yet delivers the full scenario-to-portfolio workflow. The practical solution is to pair GenAI LLMs with End-to-End Legacy Platforms such as Aladdin, BarraOne, PORT or Axioma. |
| **Workflow** | General-purpose models cannot manage trading, compliance, accounting and reporting across a single operating system. They are not the right tool for this use case. | AI is being added to incumbent platforms such as Aladdin Copilot, but the value is incremental productivity improvement rather than a fundamental leapfrogging of legacy processes. |
| **Data infrastructure** | GenAI LLMs cannot connect to GP portals, run ongoing ingestion processes, or produce audit-grade outputs across thousands of documents. The constraint is integration, not model intelligence. | Canoe Intelligence, Arch, Carta LP Portfolio Analytics and Allvue address a mechanically critical problem beyond standalone LLM capability: collecting, extracting, reconciling and distributing manager data at scale. The ROI is immediate for allocators with a meaningful manual manager data gathering process. |

**Key:** Strong Average Weak

**Source:** TNI Analysis

## Appendix

### Vantager (now MSCI Diligence Platform)

*https://www.vantager.com/*

| Brief Description & Key Product Features | Customer Segments | Value Propositions | More Information about the Company |
|---|---|---|---|
| Vantager is an AI-native pre-trade diligence platform for LPs that ingests entire GP data rooms, structures unstructured documents into standardised datasets, and surfaces inconsistencies across decks, DDQs, LPAs and track records.<br><br>**Key Features:**<br>1. **Full data room structuring:** Ingests up to 900–1,200 documents per data room via direct upload, email pipe or database connection; preprocesses, classifies and tags content at page level (LPA, PPM, track record, etc.) within 20–30 minutes.<br>2. **Diligence intelligence layer:** Auto-extracts ~75 line items today (~270+ MSCI metrics on the roadmap) across investment, operational and legal diligence; flags missing data and out-of-market terms (e.g. inconsistent hard caps, absent key person clauses).<br>3. **Reporting and comparison engine:** Generates customisable executive summaries (2–3 page IC-ready outputs), dynamic reports that re-run as new docs are added, side-by-side comparison of up to 5 strategies, and three-tier conversational AI (document, data room, full database). | 1. Endowments, foundations and pensions (manager selection, fund DD, IC-memo automation).<br>2. OCIOs, fund-of-funds and consultants (cross-manager comparison and exposure analysis at scale).<br>3. Single and multi-family offices (centralised pre-trade workflow across PE, VC, infra, hedge funds and co-investments). | 1. **Purpose-built end-to-end LP workflow:** covers the full pre-trade stack — ingestion, structuring, diligence, reporting and comparison — in a single platform, marketed as a one-stop shop for LP pre-investment activity.<br>2. **MSCI data and benchmarking integration:** post-acquisition access to MSCI/Burgiss private-markets datasets and 250,000+ indexes enables fact-checking of GP-reported numbers and a roadmap to portfolio-company-level revenue and EBITDA benchmarking.<br>3. **Institutional traction at speed:** within ~8 months of go-to-market, claimed 4,000+ funds and 100,000+ underlying investments processed across users representing $100B+ in AUM. | Founded in 2023.<br><br>HQ: New York City, US.<br><br>Latest Fundraise: Acquired by MSCI Inc. (NYSE: MSCI) on 2 March 2026; terms undisclosed. Pre-acquisition backers included BoxOne Ventures and Entrepreneur First.<br><br>**Key Personnel:**<br>1. **Mason Lender, Co-Founder and CEO:** statistician background, prior work building predictive models to assess VC manager performance. |

**Hebbia**

*https://www.hebbia.com/*

| Brief Description & Key Product Features | Customer Segments | Value Propositions | More Information about the Company |
|---|---|---|---|
| Hebbia provides an enterprise AI workspace for finance and legal teams that executes complex, document-heavy workflows with source-linked traceability.<br><br>**Key Features:**<br>1. **Multi-format data reasoning:** Works across PDFs, spreadsheets, slide decks, contracts, filings, transcripts and VDRs, with a market-leading context window as well as source-citation tracing.<br><br>2. **Third party database and cloud links:** Hebbia can link to our Preqin, Pitchbook, Snowflake, Capital IQ and FactSet.<br><br>3. **Workflow automation:** Executes repeatable workflows such as diligence Q&A, portfolio monitoring, IC memo drafting, credit-agreement parsing and reporting, with auditability of each AI action.<br><br>4. **Regulated enterprise control layer:** Built for compliance environments with granular transparency, permissioning and security controls. | 1. Asset managers, PE and private credit firms (M&A diligence, portfolio monitoring, earnings analysis, credit agreements, reporting).<br><br>2. Investment banks and advisory firms (deal-point extraction, comps and precedent transactions, IC memo and IM drafting).<br><br>3. Law firms and legal departments (contract analysis, M&A clause benchmarking, LPA side-letter tracking). | 1. **Strong cross-document reasoning:** the cleanest public example of an AI engine designed to reason over very large, mixed-format document sets and surface inconsistencies.<br><br>2. **Audit-grade transparency:** every AI output includes line-level citations and full action histories, supporting compliance and IC-grade defensibility.<br><br>3. **Marquee finance adoption:** trusted by demanding regulated enterprises; named clients include KKR, Centerview Partners and MetLife. | Founded in 2020.<br><br>HQ: New York City, US.<br><br>Latest Fundraise: US$130 million Series B led by Andreessen Horowitz with participation from Index Ventures, GV and Peter Thiel (Jul 2024); total capital raised ~US$161 million.<br><br>**Key Personnel:**<br>1. **George Sivulka, Founder and CEO**: founded Hebbia while a Stanford PhD student. |

**Clade**

*https://www.clade.co/*

| Brief Description & Key Product Features | Customer Segments | Value Propositions | More Information about the Company |
|---|---|---|---|
| Clade is an AI-native platform built for the LP/allocator side of the manager diligence workflow. It ingests an allocator's document set and organises it into a queryable, source-cited system of record, supporting diligence, fund comparison, IC-memo drafting and portfolio monitoring.<br><br>**Key Features:**<br>1. **Full-corpus ingestion and system of record:** ingests historical and ongoing decks, fact sheets, PPMs, LPAs, side letters, emails and attachments via API, manual upload or email forwarding; classifies, indexes and tags documents and sorts them into per-fund workspaces.<br><br>2. **Diligence, extraction and reporting**: side-by-side fund comparison across large document sets; an LPA matrix that extracts structured data points; IC memos, tear sheets and ODD reports generated in the allocator's own format; natural-language chat across the corpus<br><br>3. **AI harness, agents and skills**: routing across several LLMs with Python-based calculation agents; three skill types (chat, report writing and matrix extraction), with agents in development for intake, portfolio monitoring, reconciliation and natural-language fund screening. | 1. Endowments and foundations (e.g. Purdue Research Foundation, University of Florida) — manager diligence, IC-memo automation and portfolio monitoring.<br>2. Public pension funds and healthcare plans (e.g. SCERS) — enterprise deployment under demanding compliance requirements.<br>3. OCIOs, investment consultants, RIAs and family offices (e.g. Monticello Associates) — centralised, cross-manager diligence and reporting. | 1. **LP-focused design**: built for the allocator side rather than GPs, supporting intake, diligence, fund comparison, memo generation and portfolio monitoring within a single platform organised as per-fund workspaces.<br><br>2. **Workflow automation**: reusable agentic "skills" automate recurring tasks such as peer comparisons, quarterly updates and re-up diligence.<br><br>**3. Cost management**: routes across several LLMs and uses Python-based calculation agents, with usage controls intended to keep costs predictable. | Founded in 2018.<br><br>HQ: New York City, US.<br><br>Latest Fundraise: not disclosed; understood to be backed by family offices and a private equity investor.<br><br>**Key Personnel:**<br>1. **Jonathan Lipton, Founder and CEO:** previously led ultra-high-net-worth strategies at Credit Suisse. |

**Canoe Intelligence**

*https://canoeintelligence.com/*

| Brief Description & Key Product Features | Customer Segments | Value Propositions | More Information about the Company |
|---|---|---|---|
| Canoe is an AI-native data infrastructure platform for alternatives that automates the collection, categorisation, extraction and validation of post-investment documents, delivering structured private-markets data into clients' downstream analytics, accounting and reporting systems.<br><br>**Key Features:**<br>1. **Automated document collection and processing:** RPA-driven scraping of portals, plus ingestion via email and direct upload; auto-categorises and renames capital calls, distributions, K-1s, financials, quarterly reports and factsheets across PE, VC, hedge funds and real estate.<br><br>2. **Proprietary AI extraction trained on private markets:** in-house LLMs trained on the largest private-markets document set (44,000+ unique funds) extract structured data — including portfolio-company look-through (revenue, EBITDA, net debt, geography, sector) — with audit trails to source.<br><br>3. **System-agnostic data delivery:** open API to 300+ downstream integrations; recently launched Bloomberg integration positioning Canoe as the private-markets data layer alongside Bloomberg's public-markets stack. | 1. Endowments, foundations and pension funds (post-investment alts data operations and look-through reporting).<br><br>2. OCIOs, family offices and wealth managers (centralised alts data ingestion and downstream reporting).<br><br>3. Asset servicers and fund administrators (e.g. State Street, BNY Mellon, JP Morgan, Gen II) processing allocations at scale through Canoe. | 1. **Operational scale and efficiency:** clients report 15–20x increase in funds processed per employee and 5–8x faster reporting cycles.<br><br>2. **Scaled proprietary alternatives data corpus**: Canoe's fund, document and data-point coverage supports shared intelligence, extraction accuracy and emerging market-level analytics.<br><br>3. **Marquee institutional adoption:** 400+ clients including 12 of the top 30 US endowments, plus Cambridge Associates, Hamilton Lane, Blackstone, and Prime Buchholz. | Founded in 2013 (commercial launch 2018).<br><br>HQ: New York City, US.<br><br>Latest Fundraise: US$36 million Series C led by Goldman Sachs Alternatives (Jul 2024), with F-Prime Capital and Eight Roads Ventures; total capital raised ~US$72.8 million.<br><br>**Key Personnel:**<br>1. **Jason Eiswerth, CEO (since Sep 2020):** ex-Nima Capital, TheMarkets.com, Goldman Sachs and Lehman Brothers.<br>2. **Seth Brotman, Co-Founder, President and COO:** founding CEO; ex-Novus Partners. |

## True North Institute

The **True North Institute** is a research-driven investment platform focused on identifying the technologies, sectors and macroeconomic shifts most likely to shape the future of the global economy and energy transition. Combining deep fundamental analysis with a strategic institutional-investor lens, TNI conducts rigorous research into areas such as geothermal energy, long-duration storage, hydrogen, carbon markets, grid infrastructure and broader institutional portfolio strategy. The firm's work is grounded in the belief that the world's largest investment opportunities emerge where technological change, policy, capital flows and long-term economic fundamentals intersect — and that objective, evidence-based analysis is essential to identifying those opportunities early.

TNI also operates as an active investment platform, deploying Miranda family capital into companies and projects accelerating the global energy transition. In 2025, TNI co-founded the All Aboard Coalition and All Aboard Fund alongside Chris Anderson's Exa Ventures, bringing together many of the world's leading climate-tech venture and growth investors to help breakthrough energy technologies scale beyond the commercial "valley of death." Through both its research and investment activities, TNI seeks to mobilise institutional capital toward technologies capable of delivering long-term economic value while materially decarbonising the real economy.

TNI operates from its office in London, and is staffed with a small team including Stan Miranda, Jack Haynes, Tristan Varakuta and Rosanna Quayle

**Stan Miranda**
*Founder and Chief Executive Officer*

**Tristan Varakuta**
*Investment Associate*

**Jack Haynes**
*Partner, Head of Investments*

**Rosanna Quayle**
*Chief of Staff*

# Disclaimer

The information provided in this document is for informational purposes only and does not constitute a solicitation, offer, or sale of securities. Neither the investment examples cited nor mention of examples constitute investment advice or a recommendation to purchase or sell any securities.

Cautionary note: TNI has not directly tested each tool referenced. Views are based on vendor discussions and AI-assisted research.