

Sitrep: The Intelligence Layer & Trade Proposal Engine for Investment Decisions

Agents do the research and propose the trades — portfolio managers make the call

Sitrep AI

September 2026

Contact: research@sitrep.so | sitrep.so

Abstract

Active management is losing ground to passive allocation not because portfolio managers lack skill, but because the research required to sustain genuine edge has become too labor-intensive to run at scale. Funds today either stitch together a dozen point tools — a terminal for market data, a spreadsheet plug-in for model updates, a chat assistant for quick questions — or babysit general-purpose AI agents that hallucinate financial facts and default to consensus, the one output an active manager cannot use. We present Sitrep, an agentic research and trade-proposal engine built for portfolio managers, hedge funds, family offices, and RIAs. Sitrep ingests a fund's portfolio, cost basis, and mandate directly from its custodian or order management system, and ingests the fund's own proprietary research archive — its models, memos, and notes — as grounding context no generic tool has access to. Every morning before the market opens, Sitrep delivers structured trade proposals, tailored to that fund's specific book and mandate, with every material claim cited back to a filing, transcript, or live data point. Sitrep does not place trades. It is built on a simple division of labor: AI can't trade — portfolio managers can. Sitrep's role is to commoditize the research grind so that portfolio managers can spend their time on the judgment calls only they can make.

1. Introduction and Problem Statement

1.1 The Active Management Problem

Passive allocation has captured a growing share of global assets under management over the past two decades, and the reason is measurable rather than ideological: on a net-of-fees basis, most active equity managers underperform their own benchmark over long holding periods. S&P's SPIVA scorecards, the longest-running independent measure of this gap, have shown this pattern consistently across market cycles and geographies.⁵ A passive product is difficult to beat by a manager who spends the majority of the week on data assembly rather than differentiated judgment.

1.2 Why General-Purpose AI Hasn't Fixed This

The obvious response is to point AI at the problem, and funds have tried. In practice, this produces one of two outcomes. Either a fund stitches together a dozen point tools — a terminal for market data, a spreadsheet plug-in to refresh a model, a general chat assistant for quick lookups — none of which know the fund's mandate, its current book, or its house view. Or a fund deploys a general-purpose AI agent and spends the day babysitting it: verifying every number it produces, because general-purpose models hallucinate financial facts with the same fluent confidence as correct ones, and correcting its tendency to converge on consensus, which is close to worthless to an active manager who is paid specifically to disagree with the market.

1.3 The Limits of the Legacy Terminal

The terminal era solved a real problem: it put real-time market data, news, and analytics in front of every trading desk in one place. Bloomberg alone now serves more than 300,000 subscribers worldwide, at roughly \$32,000 per seat per year for a single-terminal subscription.² But a terminal is a data surface, not a research agent. It does not know a fund's mandate, has never seen its book, and cannot read the fund's own proprietary research archive — the models, memos, and house views that are the actual source of a manager's edge. Layering a data-automation tool such as Daloopa on top, which structures and source-links fundamental data extracted from filings, closes part of the gap for equity research specifically,³ but neither product was built to turn a fund's own context into a trade proposal tailored to its book.

2. Market Opportunity

2.1 The Cost of the Terminal Era

The terminal era's pricing reflects its architecture: a single, all-encompassing product, licensed per seat, with no way to pay only for the research a fund actually consumes. Bloomberg's pricing has held near \$24,000–\$32,000 per seat per year for over a decade, with modest annual increases and no meaningfully discounted tier for lighter usage;² a five-person research team pays for five full licenses regardless of how much of the terminal any one analyst actually uses on a given day.

Segment	Data Point	Source
Bloomberg Terminal	~300,000+ subscribers worldwide; ~\$24,000–\$32,000 per seat / year	Bloomberg L.P. public reporting
Daloopa structured dataset	6,000+ public companies; >99% reported data-point accuracy	Daloopa public materials
Grounded vs. ungrounded agent accuracy	~90% on structured, source-linked data vs. ~19–20% on public web sources	Daloopa FinRetrieval benchmark
Active manager underperformance	Majority of active equity managers trail their benchmark over 10+ year horizons	S&P SPIVA scorecards

Table 1: Selected reference points from the legacy terminal and data-automation layers.^{2,3,5}

2.2 The Data Automation Layer Is Already Being Built

A parallel wave of infrastructure has emerged to solve a narrower problem: getting clean, source-linked financial data out of filings and into a model or an AI agent reliably. Daloopa, for example, structures financial data from SEC filings and earnings materials across more than 6,000 public companies, hyperlinking every extracted figure back to its original source and reporting an accuracy rate above 99% across the data points it structures.³ Independent benchmarking of agents connected to this kind of structured data illustrates why grounding matters: agents working against structured, source-linked data have been measured at roughly 90% task accuracy, versus roughly 19–20% when the same agents work directly against public web sources.³ This is a necessary foundation for reliable financial AI — but a clean data layer is not, by itself, a research agent that knows a specific fund's mandate and book.

2.3 Why Now

Three developments converge to make an agentic research layer viable for the first time. First, structured, source-linked financial datasets of the kind Daloopa and similar providers now offer give AI agents something reliable to ground claims in. Second, institutional data plumbing — custodian and order-management-system APIs, portfolio and cost-basis feeds — has matured to the point where a fund's live positions can be ingested directly rather than re-keyed by hand. Third, agentic AI frameworks have matured past single-shot question answering into genuine multi-step research and proposal generation. What has been missing is a product that combines all three around a specific fund's own mandate and book, rather than treating every user identically.

3. The Sitrep Thesis: AI Can't Trade. Portfolio Managers Can.

We view the automation of investment research as having progressed through four phases.

Phase I — Manual Research. Analysts read filings, build models by hand, and maintain coverage the old-fashioned way: bound by the number of hours in a day and the number of tickers one person can reasonably track.

Phase II — Point Tools and Terminals. Bloomberg-style terminals, spreadsheet plug-ins, and data-automation layers such as Daloopa each solve a piece of the problem, but a human still has to stitch data, model, and mandate into a single view every single day.

Phase III — General-Purpose AI Agents. Flexible and fast, but prone to hallucinated financial facts and consensus-shaped output — a fund still has to babysit and correct them, which erodes much of the time saved.

Phase IV — Mandate-Aware Agentic Research (Sitrep). An agent that knows a specific fund's mandate, book, and proprietary research context, grounds every claim to a source, and proposes trades — but stops short of executing them.

AI can't trade. Portfolio managers can.

Sitrep's thesis is that the industry does not need an AI that occasionally produces a good idea. It needs an AI that reliably compresses the research grind — data assembly, model maintenance, first-pass diligence — into minutes, so that the portfolio manager's day is spent on the judgment calls that a mandate, a book, and a career are actually built on.

4. System Architecture

Sitrep is built around five architectural commitments, summarized in Table 2 and described below.

Layer	Function
1 — Ingestion (the Vault)	Ingests portfolio, cost basis, mandate, and the fund's own proprietary research archive
2 — Grounding and Citation	Every material claim is linked to a specific, checkable source
3 — Research and Proposal Engine	Produces research memos, dashboards, and morning trade proposals
4 — Disagreement Preservation	Actively resists converging on consensus views
5 — Execution Boundary	Delivers proposals only; has no pathway to place a trade

Table 2: The five layers of Sitrep's architecture.

4.1 Ingestion Layer (The Vault)

Sitrep connects to a fund's custodian, order management system, or accepts direct upload, and ingests positions, P&L, and cost basis — every subsequent answer is grounded in what the fund

actually owns, not a generic universe of tickers. Alongside portfolio data, Sitrep ingests the fund's own proprietary context: its investment mandate, thesis documents, sell-side research it subscribes to, and its internal research archive of models, memos, and notes. No terminal has access to this material, because it is not public; it is the specific, idiosyncratic body of work that makes one fund's edge different from another's, and it is the primary grounding context behind every signal, backtest, and trade proposal Sitrep produces.

4.2 Grounding and Citation Layer

Every material financial claim Sitrep produces is linked to a specific source — a filing, an earnings transcript, an options chain, a live price feed, or a document the fund itself uploaded. An unsourced claim is treated as a defect, not an acceptable trade-off for speed. This is the same discipline that makes a source-linked data layer valuable in isolation; Sitrep applies it not just to raw data points but to every downstream research output and trade proposal built on top of them.

4.3 Research and Proposal Engine

Sitrep's research agents run against a fund's ingested context to produce structured outputs: research memos, cross-asset dashboards, and, every morning before the market opens, a trade proposal tailored to that fund's specific book and mandate. A trade proposal is not a generic idea; it is sized and framed against positions the fund already holds, cited to the specific facts that motivate it, and timed to be actionable before the trading day begins.

4.4 Disagreement Preservation

Sitrep is explicitly built not to converge on consensus. An active manager who is charged fees to outperform the market's average view has no use for a tool whose default behavior is to reproduce that view. Sitrep's research agents are configured to surface genuinely independent, differentiated conclusions — and to disagree with each other, and with the market, when the evidence supports it — because consensus is the one output an active manager cannot monetize. Section 5.1 formalizes why this matters.

4.5 The Execution Boundary

Sitrep is not an auto-trader, and it does not sit between a portfolio manager and a broker's execution API. Every trade proposal lands in a portfolio manager's hands as a proposal, not an instruction; the decision to act, size, and execute remains entirely with the human who is accountable for it. This boundary is architectural, not a policy choice that could be quietly relaxed — Sitrep has no execution pathway to remove.

5. Mathematical Foundations

This section formalizes two properties underlying Sitrep's design. Mathematical notation is used only where it clarifies a claim made elsewhere in this paper.

5.1 The Diversity Prediction Theorem

Sitrep's disagreement-preservation design (Section 4.4) is not merely a stylistic preference; it has a formal basis in the mathematics of collective forecasting. For a set of independent forecasts of a single quantity, the error of their average is related to the average individual error and the diversity among the forecasts by:

$$\text{Collective Error} = \text{Average Individual Error} - \text{Prediction Diversity}$$

This identity, first derived in the ensemble-learning literature⁶ and later popularized in the context of group decision-making,⁷ shows that a group of forecasters that converges toward a single consensus view mechanically drives diversity toward zero, and with it, any advantage the average forecast might have held over the typical individual forecast. Applied to investment research: a tool that nudges every user toward the same conclusion does not just fail to add value — it actively erodes the main statistical advantage an independent, differentiated view can offer.

5.2 Grounding as a Formal Constraint

Let a research claim c be admissible only if it can be mapped to a verifiable source s via a grounding function g , such that $g(c) = s$ and s is independently checkable by the portfolio manager receiving c :

$$g(c) = s, \text{ where } s \text{ is independently verifiable}$$

Sitrep enforces this constraint on generation rather than as a post-hoc filter: a claim without a valid $g(c)$ is not produced in the first place, rather than being generated and then flagged afterward. This mirrors, in the investment-research domain, an architectural discipline increasingly used across AI applications in regulated or high-stakes industries: separating the parts of a system that must be flexible from the parts that must be verifiable, rather than asking one generative process to be both.

6. Platform Modules

Sitrep exposes its research engine through five interoperable modules.

6.1 Vault

Vault is Sitrep's ingestion layer: it connects to a fund's custodian or OMS for live portfolio, P&L, and cost-basis data, and ingests the fund's own mandate, thesis documents, sell-side research, and internal research archive. Vault is the grounding context every other module draws on.

6.2 Research Agent

Research Agent answers ad hoc questions and produces on-demand analysis — comps, scenario models, sector comparisons — grounded in Vault's context and cited to source at every material claim.

6.3 Morning Sitrep

Morning Sitrep is the daily trade-proposal briefing: before the market opens, it delivers structured, cited trade proposals sized against the fund's actual book and framed against its mandate — ready to act on, not a generic idea list.

6.4 Dashboards

Dashboards are AI-generated widgets on a flexible canvas that maintain themselves indefinitely — a cross-asset correlation view, a factor-exposure summary, a watchlist — refreshed automatically against live data rather than rebuilt by hand each morning.

6.5 Sitrep Mail

Sitrep Mail lets a portfolio manager forward any email to Sitrep and receive a reply in the same thread, grounded in the fund's portfolio, Vault context, and live market data — no separate interface required for a quick question.

7. Benchmarks and Grounding Evidence

Sitrep's grounding discipline (Section 4.2) is not a theoretical preference; independent, third-party benchmarking of the broader category gives a sense of what is at stake. In Daloopa's published FinRetrieval benchmark, agents connected via MCP to structured, source-linked financial data achieved approximately 90% task accuracy, compared with roughly 19–20% accuracy for comparable agents working directly against public web sources.³ This is a third-party result describing the broader category of grounded-versus-ungrounded financial agents; it is not a claim about Sitrep's own measured performance, and we cite it here as the clearest available evidence for why the architecture described in Sections 4 and 5 is built the way it is.

Sitrep's own benchmark disclosures — accuracy, latency, and proposal-quality metrics specific to its research and proposal engine — are governed by the same grounding discipline described in Sections 4.2 and 5.2, and are held to the same third-party standard of verifiability the architecture is built around.

8. Security and Trust Architecture

Encryption. Fund data is encrypted at rest using AES-256 and in transit using TLS 1.2 or higher.

Independent Audit. Sitrep is pursuing a SOC 2 Type II audit as part of its enterprise readiness program.

No Execution Pathway. Sitrep is not a broker-dealer and does not place trades. Every output is a proposal delivered to a human decision-maker, consistent with the architectural boundary described in Section 4.5.

Tenant Isolation. Each fund's Vault — its portfolio, mandate, and proprietary research archive — is walled off from every other fund's data and is not pooled across customers without explicit consent.

Grounded by Default. As described in Sections 4.2 and 5.2, an unsourced claim is not generated in the first place — hallucinations are treated as a defect in the architecture, not a risk to be disclosed and tolerated.

9. Institutional Use Cases

10.1 Hedge Funds

Hedge funds use Sitrep to compress the daily research grind — idea generation, thesis checks, scenario modeling — into a morning proposal tailored to their book, freeing portfolio managers and analysts to spend the trading day on judgment calls rather than data assembly.

10.2 Family Offices

Family offices, often running lean teams across a wide mandate, use Sitrep to maintain coverage across a broader set of positions and asset classes than a small internal team could otherwise track continuously.

10.3 RIAs

Registered investment advisors use Sitrep to support client portfolio reviews with grounded, cited research and proposals, without maintaining the size of research team a Bloomberg-plus-analyst model traditionally required.

10. Business Model

Sitrep prices consumption, not seats. A quick price check and a full, mandate-aware trade proposal are not the same amount of work, and a seat-based license charges identically for both regardless of how a fund actually uses the product.

Tier	Example	Pricing Unit
Quick Check	Spot price, single-metric lookup	Low, per query
Research Brief	Multi-source thesis check, sector or comp comparison	Moderate, per brief
Trade Proposal	Full, mandate-aware, book-tailored proposal with citations	Highest, per proposal

Table 3: Sitrep consumption-based pricing tiers.

This model scales naturally with fund size and usage intensity, in contrast to the seat-based terminal model described in Section 2.1, where a five-person team pays for five full licenses regardless of how lightly some of those seats are actually used.

11. Competitive Landscape

Sitrep sits in a landscape with three adjacent categories, summarized in Table 4.

Category	Examples	Strength	Gap Sitrep Fills
Legacy terminals	Bloomberg, FactSet, Refinitiv/LSEG	Comprehensive real-time data, news, and messaging in one place	Seat-based, generic to every user; does not know a fund's mandate, book, or proprietary research

Data automation layer	Daloopa, AlphaSense	Structured, source-linked fundamental data; strong model-update automation	A data and modeling layer, not a mandate-aware research or trade-proposal agent
Closest peer	Kimpton AI	AI-native research platform for portfolio managers; publicly reports tracking over \$3.3B in assets on its platform, backed by Y Combinator	Validates the category; Sitrep is a separate entrant building the same class of mandate-aware, disagreement-preserving research agent

Table 4: Competitive landscape across legacy terminals, data automation, and mandate-aware research agents.^{1,2,3}

None of the legacy terminal vendors were built as agentic research platforms, and the data automation layer — while an excellent foundation — stops at structured data rather than a proposal grounded in a specific fund's mandate and book. Kimpton AI's public traction is the clearest evidence that institutional demand for this category is real;¹ Sitrep is built around the same core thesis — AI does the research, the portfolio manager makes the trade — as an independent entrant into the same emerging category.

12. Vision and Roadmap

Investment research is becoming an agentic workflow rather than a human-hours-bound one. Table 5 summarizes Sitrep's roadmap.

Phase	Horizon	Key Milestones
I — Core Platform	2026 H2	Vault, Research Agent, and Morning Sitrep live; custodian and OMS integrations for portfolio and cost-basis ingestion; consumption-based pricing in production.
II — Platform Expansion	2027 H1	Dashboards and Sitrep Mail generally available; expanded custodian and OMS integrations; published benchmark disclosures for Sitrep's own research and proposal accuracy.
III — Asset Class Expansion	2027 H2	Extend beyond public equities into fixed income and derivatives research; deeper sector and factor-model coverage.
IV — Category Scale	2028+	Broaden platform adoption across hedge funds, family offices, and RIAs; expand internationally alongside institutional data-access partnerships.

Table 5: Sitrep roadmap from initial rollout to category scale.

Over a longer horizon, as Sitrep's grounding and disagreement-preservation discipline extends across more asset classes and data sources, the objective is the same one that motivates the whole platform: let portfolio managers spend their time on the decisions only they can make, not the research required to reach them.

13. Conclusion

The next generation of investment research infrastructure will not be another terminal, and it will not be a general-purpose AI agent that has to be babysat and corrected all day. It will be an agent that knows a fund's mandate and book as well as its own analysts do, grounds every claim to a source, and resists the pull toward consensus that active managers are paid specifically to avoid.

Sitrep's thesis is that the bottleneck in active management is not talent or conviction — it is the sheer amount of ungrounded, unmandated, generic research work standing between a portfolio manager and the decision only they can make. By ingesting a fund's own portfolio and proprietary context, grounding every claim to source, preserving genuine disagreement, and stopping short of execution, Sitrep aims to become the research layer the buy side runs on.

AI can't trade. Portfolio managers can. Sitrep exists to make sure they spend their time doing exactly that.

Glossary

Cost Basis	The original value of a security for tax and P&L purposes, used to calculate gains or losses on a position.
Custodian	A financial institution that holds a fund's securities and other assets in safekeeping, and is typically the source of a fund's official position and cost-basis records.
Diversity Prediction Theorem	A formal result from ensemble forecasting showing that the error of a group's average forecast equals the average individual error minus the diversity among the forecasts.
Family Office	A private wealth-management firm that manages investments and other affairs for a single wealthy family or a small group of families.
Grounding	The practice of linking every generated claim to a specific, independently verifiable source, such as a filing, transcript, or live data feed.
Mandate	The set of objectives, constraints, and permitted strategies that govern how a fund is allowed to invest.
MCP	Model Context Protocol. An open standard enabling AI agents to invoke external tools and data sources through a standardized interface.
OMS	Order Management System. Software used by trading desks to manage the lifecycle of orders from creation through execution and settlement.
Passive Allocation	An investment approach that seeks to replicate a market index rather than actively select securities expected to outperform it.
RIA	Registered Investment Advisor. A firm or individual registered with regulators to provide investment advice to clients, typically for a fee.

SOC 2 Type II	An independent audit standard assessing whether a service organization's controls over security, availability, and confidentiality operated effectively over a period of time.
SPIVA	S&P Indices Versus Active. A long-running scorecard series published by S&P Dow Jones Indices comparing active fund performance against benchmark indices.
Trade Proposal	A structured, sourced recommendation to buy, sell, or adjust a position, sized and framed against a fund's actual holdings and mandate, but not an executed instruction.

References

- [1] Kimpton AI (2026). Company Profile and Product Documentation. Y Combinator; kimpton.ai.
- [2] Bloomberg Terminal (2026). Subscriber count and pricing data. Bloomberg L.P.; Wikipedia; industry pricing surveys (Vendr, CostBench).
- [3] Daloopa (2026). Product Documentation, FinRetrieval Benchmark, and Press Releases. daloopa.com.
- [4] Anthropic (2024). Model Context Protocol Specification. <https://modelcontextprotocol.io>
- [5] S&P Dow Jones Indices (2026). SPIVA U.S. Scorecard.
- [6] Krogh, A. and Vedelsby, J. (1995). Neural Network Ensembles, Cross Validation, and Active Learning. Advances in Neural Information Processing Systems (NeurIPS).
- [7] Page, S. E. (2007). The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. Princeton University Press.

This whitepaper is provided for informational purposes only and does not constitute financial, legal, investment, or tax advice. Forward-looking statements involve risks and uncertainties and should not be relied upon as guarantees of future performance. Sitrep is not a broker-dealer or registered investment advisor and does not place trades on behalf of any client. References to third-party products (Bloomberg, Daloopa, Kimpton AI, and others) describe publicly available information about those products as of the date of writing and do not imply endorsement by, or affiliation with, those companies. © 2026 Sitrep AI.