

Investment Research Agents and Protection of Investment Context

An Asset and Wealth Executive Decision on Confidential Knowledge

Phil Hatch | Akholi

October 23, 2026

Investment executives should govern the confidential context that provider agents receive, infer, and retain. Portfolio intentions, client circumstances, research priorities, and investment rationale can require protection even when individual source documents are publicly available.

Identify the Investment Context the Service Creates

The executive should ask what the provider can learn from the sequence of research requests, holdings, watchlists, client information, and decisions. Public market facts combined with private instructions can reveal a strategy or customer circumstance. The protection decision should follow that context rather than only the classification of each source document.

LSEG's connector discussion describes customer-controlled tools and context in agentic workflows.¹ It is a provider product account, not a guarantee of confidentiality for every configuration. The investment firm should determine what context it exposes and retains across the service.

For a hypothetical manager evaluating a position, a set of public-data queries may reveal an intended trade or concern about a holding. The provider's logs, memory, support access, and shared analytics should be included in the assessment, even if no restricted research report was uploaded.

¹ Emily Prince, "The Hidden Economics of AI Agents—and How LSEG's MCP Connector Helps Improve Token Efficiency," LSEG, July 17, 2026, licensed content, permissioned access, and customer-controlled agent workflow, <https://www.lseg.com/en/about-us/new-at-lseg/the-hidden-economics-of-ai-agents-and-how-lsegs-mcp-connector-helps-improve-token-efficiency>.

Separate Service Use from Further Learning

The study describes operational knowledge and downstream dependencies outside a single formal record.² It does not establish investment-information leakage. The transferable question is whether the buyer knows where valuable context accumulates and which parties can use it.

Define permitted use for delivery, support, incident investigation, dedicated improvement, shared-product development, and broader training. These are different purposes. A general promise not to train a foundation model may leave other uses of logs, derived context, or workflow learning unresolved.

Classify buyer-specific information, provider background methods, third-party licensed content, and newly created artifacts. Ownership alone does not settle confidentiality, permission, access, or retention. The executive should require a usable allocation that supports the investment function and protects relevant duties.

Use an Investment Context Protection Record

Figure 1 proposes four elements: context, recipient, permitted purpose, and lifecycle control. Identify the data and inferences needed for the service, their sensitivity, authorized access, storage, retention, and deletion. Preserve a route to investigate material use.

Morgan Lewis's contracting analysis addresses data, output, and training-use provisions.³ The proposed record applies those questions to investment context. It is not a legal conclusion that every inference has the same ownership or protection in every jurisdiction.

Test representative prompts, retrieval, tool calls, memory, logs, and support cases. Determine whether unnecessary context can be withheld or reduced. Verify that a correction or deletion request reaches relevant downstream copies where required and feasible, with limitations disclosed.

² Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 4, "Knowledge at Risk of Departure," <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 3, Third-Party Ecosystem, sec. 6.1, "Mapping of Third and Fourth Party Dependencies," indicator C196.

³ Doneld Shelkey, "Negotiating AI Provisions in Commercial and Technology Contracts: Where the Market Is Heading," *The Licensing Journal*, May 1, 2026, pp. 1–2, intellectual property and transition, <https://www.morganlewis.com/-/media/files/publication/outside-publication/article/2026/negotiating-ai-provisions-in-commercial-and-technology-contracts-where-the-market-is-heading-the-licensing-journal.pdf>.

Figure 1. The investment context protection record

Protection layer	Required evidence	Executive question
Context	Strategy client and inferred information	What valuable knowledge is exposed
Recipient	Provider model tools support and downstream use	Who can receive or infer it
Purpose	Delivery dedicated improvement and further reuse	What uses are permitted
Lifecycle	Access retention investigation and transition	Can the firm maintain practical control

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Preserve Local Duties and Practical Control

EU, UK, Swiss, US, and other investment and wealth activities can involve different confidentiality, privacy, market-conduct, recordkeeping, and client duties. Qualified specialists should establish applicable conclusions. A global research workflow should preserve material legal-entity and client distinctions.

The provider should disclose material recipients, locations, and changes in use. Contract terms should support evidence access, restriction, incident cooperation, and transition. Confidentiality protections should be assessed against actual support and downstream operating practices rather than a prime-level assurance alone.

Retain enough buyer expertise and context to challenge the research and continue the function. Protecting information should not leave the firm unable to understand its own decision history after exit. Balance minimization with the records needed for authorized review and applicable retention.

Approve Context Exposure as an Investment Decision

A reasonable objection is that useful research requires rich context. The choice is not between unrestricted disclosure and no context. Define the minimum suitable information and approved purposes for each task, with stronger controls where the consequence warrants them.

Before material use, require a context-flow map, permitted-use conclusions, tested controls, unresolved limitations, and an incident route. Review changes to models, memory, tools, recipients, and client scope. Restrict uses whose protection cannot be substantiated.

The executive should approve an agentic research service when the firm understands what investment context it exposes and how that context remains governed. Better analysis can create value, but the delivery model should preserve the confidentiality, judgment, and continuity on which the investment business depends.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁴

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is an investment-context protection record addressing confidential inferences and accumulated workflow knowledge as well as individual source documents. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

- Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.
- Prince, Emily. “The Hidden Economics of AI Agents—and How LSEG’s MCP Connector Helps Improve Token Efficiency.” LSEG. July 17, 2026. <https://www.lseg.com/en/about-us/new-at-lseg/the-hidden-economics-of-ai-agents-and-how-lsegs-mcp-connector-helps-improve-token-efficiency>.
- Shelkey, Doneld. “Negotiating AI Provisions in Commercial and Technology Contracts: Where the Market Is Heading.” *The Licensing Journal*. May 1, 2026. <https://www.morganlewis.com/-/media/files/publication/outside-publication/article/2026/negotiating-ai-provisions-in-commercial-and-technology-contracts-where-the-market-is-heading-the-licensing-journal.pdf>.

⁴ Hatch, *Agentic AI Outsourcing*, Introduction, “Executive Summary” and “Evidentiary Basis and Its Limits.”