BUILT BY A TEAM THAT HAS SECURED MAJOR INSTITUTIONS WORLDWIDE FOR 25 YEARS.

ARGENT SESSION REPLAY

REPLAY WHAT YOUR AGENT SENT AND GOT BACK. PROVE WHO READ IT.

Argent Strategic Intelligence puts a checkpoint between your AI agents and the AI models they call. Unless an admin switches it off, it keeps what each agent sent and what came back, every request in order, encrypted per tenant. When a customer complains or an examiner asks, you open the conversation itself and step through it.

Reading the words is itself on the record. The reader confirms the read is recorded under their name, and Argent writes that signed record before any text appears. No record, no text. Only the people responsible for an agent and your top admins can read it; an auditor sees that a prompt exists and its fingerprint, not the words.



SEE.
Every request in a conversation, in order.



PROVE.
A signed record of who read each prompt.



STOP.
No signed record of the read, no text shown.

FIVE THINGS TO KNOW ABOUT SESSION REPLAY

1	**Every conversation** Which letter did the agent draft for her?	Argent groups each agent's requests into conversations. The Sessions page lists them with owner, client, models, turns and refusals. Open one and its timeline shows every request in order, subagent turns indented under the request that started them.
2	**Step through it** What happened to the request that drafted it?	Open a request and press Play: the call arrives, goes to the model, and the answer returns. Each step shows its decision and the signed row behind it. Replay walks through what was recorded; it does not re-run anything.
3	**The audited read** Compliance needs the exact words.	Reading stored text is a separate act. The reader ticks "This read is recorded in the signed audit log under my name." Argent writes that signed record first, then shows the prompt and the reply. If the record cannot be written, no text is shown.
4	**Who may read** An auditor is reviewing the complaint.	Only the agent's owner, the department and organization admins over that owner, the tenant admin and the SuperAdmin can read the words. An auditor or operator sees that a prompt was stored, its size and its fingerprint, never the text.
5	**You choose what is kept** Isn't storing prompts a new risk?	An admin chooses what is stored: prompts and replies, prompts only, or nothing. Stored text is encrypted per tenant and tied to the signed record by a fingerprint. How long it is kept is a setting, and older text ages out to an archive.

WHAT THE RULES ALREADY ASK

NY-regulated financial firms. 23 NYCRR 500.6 asks covered entities for "audit trails designed to detect and respond to cybersecurity events."[1]

HIPAA covered entities. 45 CFR 164.312(b) asks for mechanisms "that record and examine activity in information systems that contain or use electronic protected health information."[2]

HOW A TEAM STARTS

1. Route. Send your agents through Argent. Conversations are kept unless an admin switches it off.

2. Choose. Decide what is kept: prompts and replies, prompts only, or nothing.

3. Replay. Open Sessions, step through a conversation, and reveal its text with a signed read.

BOOK A DEMO. WE WILL REPLAY A CONVERSATION FOR YOU, WITH EVERY READ ON THE RECORD.

NYEDIS.AI/DEMO

Sources. 1 New York State Department of Financial Services, 23 NYCRR 500.6(a)(2), Audit trail (Second Amendment, 2023). 2 HIPAA Security Rule, 45 CFR 164.312(b), Audit controls (Cornell LII copy of the eCFR).