GOAL DATA SHEET: DATA LOSS PREVENTION FOR AI REQUESTS AND REPLIES

ARGENT DATA LOSS PREVENTION

YOUR DATA NOW LEAVES INSIDE AI REQUESTS. STOP IT AT THE GATE.

If your goal is data loss prevention, the leak path has moved. Customer records, account numbers and keys now leave inside AI requests, typed by a person or pulled in by an agent's tools, and models can repeat them back. In one study, 8.5% of employee prompts to AI tools held sensitive information.[1] **Argent Strategic Intelligence** sits between your AI agents and the models they call.

It finds the AI tools in use, replaces sensitive values with a redaction marker before a request leaves, screens what the model sends back and keeps real keys out of agents' hands. **Every decision becomes a signed record of what kind of data and how many, never the value. If a check cannot run, the request is refused, not sent.**



SEE.
The AI tools in use, and who uses each one.



PROVE.
A signed record of data type and count, never the value.



STOP.
Redact, withhold the reply, or refuse the request.

FIVE THINGS A DATA-LOSS PROGRAM NEEDS FROM AI

1	**Unapproved AI tools** Staff paste work into a chat tool nobody approved.	Argent reads a copy of your web proxy's access log, with nothing installed and no traffic moved. Each AI tool shows up with the login or device that used it, so you know where controls need to go first.
2	**Data in the request** A customer's SSN rides along in an agent's prompt.	The default policy replaces Social Security, card, bank account and other numbers with a redaction marker before the request leaves. Add your own patterns, or fingerprint your customer lists with Exact Data Match.
3	**Keys and passwords** A developer pastes a cloud key into a coding agent.	About 60 credential rules catch cloud, AI-provider and private keys. Agents never hold your real provider key: Argent attaches it from your vault on the way out, and a request carrying a key in its web address is refused.
4	**Data in the reply** The model repeats a customer record back.	Response screening runs the same check on what the model sends back. Redact replaces the sensitive spans; Block withholds the reply. Image and link tricks that carry data out are neutralized.
5	**Proof for audit** An auditor asks what the model was actually sent.	Replay shows what the agent sent beside what Argent forwarded. Reading the raw copy takes a deliberate reveal, recorded under the reader's name.

WHERE THE PRESSURE COMES FROM

FTC Safeguards Rule, 16 CFR 314.4. Non-bank financial institutions must limit customer information to authorized users and oversee their service providers.[2]

NYDFS, 23 NYCRR 500.17. Covered entities must notify the superintendent as promptly as possible and no later than **72** hours after determining a cybersecurity incident has occurred.[3]

HOW A TEAM GETS THERE

1. Find. Feed Argent your proxy logs and see which AI tools are in use, and from which logins or devices.

2. Watch. Argent records what it would have redacted while requests pass unchanged.

3. Enforce. Switch data protection to Enforce; the default policy redacts from the next request.

BOOK A DEMO. WE WILL SHOW IT WORKING ON YOUR OWN AI TRAFFIC.

NYEDIS.AI/DEMO

Sources. 1 Harmonic Security, "New research: the data leaking into GenAI tools", April 16, 2025 (prompts from 2024). 2 16 CFR 314.4, FTC Standards for Safeguarding Customer Information, read 2026-11-19. 3 23 NYCRR 500.17(a)(1), NYS DFS, read 2026-11-19.