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

ARGENT SENSITIVE DATA REDACTION

ONE PROMPT CAN CARRY A CUSTOMER'S SSN. STOP IT AT THE GATE.

Argent Strategic Intelligence reads each AI request routed through its gateway before anything leaves for the model. A Social Security, card or bank account number in what a person typed, or what an agent's tools brought back, is swapped for a redaction marker; the rest goes through and the agent still gets its answer. In one study, 8.5% of employee prompts to AI tools held sensitive information.[1]

Detection runs inside your own Argent deployment, and the default policy is on before anyone writes one. Every redaction becomes a signed record of what kind of data and how many, never the value itself. **If Argent cannot finish the check, the request is refused, not sent.**



SEE.
The raw request beside what the model received.



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



STOP.
Redact the value, or refuse the whole request.

FIVE THINGS A SECURITY TEAM ASKS FIRST

1	**Found without a label** The number sits in the text with no "SSN" beside it.	Argent recognizes a dashed Social Security number by its shape, and scores it higher when words like "social security" sit nearby. Card numbers and IBANs are checked against their checksums, not just pattern-matched.
2	**On by default** Nobody has written a policy yet.	With no saved policy, Argent already redacts SSN, card, bank account, IBAN, passport, driver license and medical license numbers, emails, phone numbers, IP addresses and crypto wallets, and counts secrets such as API keys.
3	**Your own formats** Customer and account numbers in your own style.	An admin writes a custom pattern, tests it on sample text and sees the matches before it goes live. Exact Data Match fingerprints your real customer records on your side; only scrambled, one-way fingerprints are uploaded, never the values.
4	**Redact, record or refuse** Some data must never leave at all.	Choose Redact + Audit, Audit Only or Deny. Deny refuses the request with a plain message that names the data type and a trace id, never the value. Masking with asterisks is available too.
5	**Proof for audit** An auditor asks what the model was actually sent.	An authorized administrator can open a request and see what the agent sent beside what Argent forwarded. Reading the raw copy takes a deliberate reveal and is itself recorded in the signed log 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 give notice no later than **72** hours after determining a cybersecurity incident has occurred.[3]

HOW A TEAM STARTS

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

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

3. Tune. Add your own patterns, switch to Deny to refuse any request with detected data, review in Replay.

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.