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

ARGENT FOR RETAILERS

AI AGENTS NOW ACT LIKE STAFF. PUT A CHECKPOINT IN FRONT OF EVERY ONE.

Argent Strategic Intelligence puts a checkpoint between your AI agents and the outside AI services they use. It checks every request before it leaves, and it runs entirely inside your own walls. The average data breach now costs $4.99 million, and one in four malicious breaches is AI-enabled.[1]

One pane of glass for your entire agent ecosystem: every agent, model, tool, owner and dollar in one view. **No rip and replace.** Argent works with what you already run: your identity provider (Okta, Microsoft Entra, LDAP), your own secrets vault (HashiCorp Vault, AWS Secrets Manager, CyberArk Conjur) and the AI tools your staff already use.

SEE.
Every agent, its owner, where it sends data.

PROVE.
A signed record of every decision.

STOP.
Unapproved models, tools and leaks are refused.

THE FIVE THINGS RETAILERS ASK FOR FIRST

1	**Budget controls** A pricing agent loops on a promotion rule all weekend.	Set spend caps by team, person and agent. An agent that hits its cap is stopped; its owner asks for more, and nothing is granted until a manager approves, in Argent or through your ServiceNow or Jira tickets. Project cost tracking shows what each project or brand spends on AI.
2	**Data protection** A support rep pastes a shopper's order and card into an AI.	Argent reads each request first and blanks out card numbers, emails and phone numbers before anything leaves (names and addresses once you switch them on, loyalty numbers once you add your own patterns). The record notes what kind of data, never the data itself.
3	**Auditing for auditors** An auditor asks which agent sent shopper data, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An outside auditor or your internal audit team can check a package of it on their own computer and see whether anything was altered.
4	**Model control** A marketing agent runs on an AI model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for customer data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Shopper data must never leave the company's own network.	Argent runs entirely inside your own company, even on a fully air-gapped network. The install bundle is checked at every step as it crosses into your closed environment.

WHAT REGULATORS ALREADY EXPECT

California Consumer Privacy Act. Fines up to **$2,663** per violation and **$7,988** per intentional violation involving California residents' data.[2]

FTC Consumer Review Rule, 16 CFR 465. Fake consumer reviews, AI-generated ones included, carry civil penalties of up to **$53,088** for each violation.[3]

HOW A RETAILER STARTS

1. Watch mode. A real-time audit and view of what your AI is doing today. Nothing is blocked.

2. Enforce mode. Switch on blocking one area at a time: shopper data, budgets, models, then tools.

3. Inside your walls. Runs on your own hardware or in your own cloud VPC. Nothing leaves the perimeter without your explicit approval.

BOOK A DEMO. WE WILL SHOW YOU ARGENT WORKING ON YOUR STORES' OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 IBM, Cost of a Data Breach Report 2026, release July 29, 2026 (all industries). 2 California Privacy Protection Agency, CCPA monetary thresholds, Cal. Civ. Code 1798.155(a), effective Jan 1, 2025. 3 16 CFR 465.2, FTC Consumer Review Rule; FTC, Dec 2025, 2025 penalty level kept for 2026 (91 FR 58446).