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

ARGENT FOR FOOD AND BEVERAGE

AI AGENTS NOW TOUCH RECIPES. 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 company. Two in five food recalls in 2025 were for an undeclared allergen,[1] and 43% of security incidents now involve AI the company never approved.[2]

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 FOOD COMPANIES ASK FOR FIRST

1	**Budget controls** A plant's scheduling agent runs up a large AI bill overnight.	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 scientist's agent sends a formula sheet to an outside AI.	Argent reads each request first and blanks out email, phone and card numbers in a consumer complaint before anything leaves the company, plus the formula codes, supplier ids and formula sheets you register with it. The record notes what kind of data, never the data itself.
3	**Auditing for auditors** The recall team asks which agent changed the label.	Every AI decision and tool call goes into a signed, tamper-evident record: which agent, which tool, on which server, and when. An auditor can check a package of it on their own computer and see whether anything was altered.
4	**Model control** A plant agent sends shift logs to a model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for formulation data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Formulas must never leave the building.	Argent runs entirely inside your 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

Reportable Food Registry, 21 U.S.C. 350f. Once a company determines a food is reasonably likely to cause serious harm or death, it must tell the FDA within **24** hours.[3]

Food Traceability Rule, 21 CFR 1.1455. From July 20, 2028, traceability records must reach the FDA "upon request, within **24** hours."[4]

HOW A FOOD COMPANY 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: formulas, budgets, models, tools.

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

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

NYEDIS.AI/DEMO

Sources. 1 U.S. PIRG Education Fund, Food for Thought 2026 (FDA and USDA recalls, 2025), Feb 2026. 2 IBM, Cost of a Data Breach Report 2026, July 2026. 3 21 U.S.C. 350f(d)(1), Reportable Food Registry, read 2026-11-14. 4 21 CFR 1.1455(c)(1); compliance date per Congressional Research Service R48925, Apr 2026.