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

ARGENT FOR INSURERS

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. Nearly nine in ten auto insurers say they use, plan to use or plan to explore AI,[1] and a state regulator found 908,547 claims violations at one insurer.[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 INSURERS ASK FOR FIRST

1	**Budget controls** A claims agent runs up a large bill after a storm.	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 line of business spends on AI.
2	**Data protection** An adjuster pastes a claimant's whole file into an AI.	Argent reads each request first and blanks out Social Security, card and bank account numbers before anything leaves, plus your own claim and policy numbers once you add their formats. The record notes what kind of data, never the data itself.
3	**Auditing for auditors** An examiner asks which agent sent what, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An examiner or internal audit can check a package of it on their own computer and see whether anything was altered.
4	**Model control** Staff start using an AI model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for claimant data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Claimant data must never leave the building.	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

NAIC Insurance Data Security Model Law. The model law gives **72** hours to tell the commissioner, and a vendor breach counts too; adopting states can set their own deadline.[3]

New York DFS, Circular Letter 7 (2024). Insurers keep documentation of the AI and outside data used in underwriting and pricing, ready on request.[4]

HOW AN INSURER 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: claimant 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 OWN CLAIMS AND PRICING USE CASES.

NYEDIS.AI/DEMO

Sources. 1 NAIC, Artificial Intelligence topic page, updated Apr 3, 2026. 2 California Department of Insurance, press release, Sept 21, 2018. 3 NAIC Insurance Data Security Model Law (#668), Sec. 6, 2017. 4 NY DFS Insurance Circular Letter No. 7 (2024), July 11, 2024.