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

ARGENT FOR HOTELS AND RESORTS

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. A hospitality data breach now costs $4.33 million on average, and across industries 43% of security incidents involve AI the company never approved.[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 HOTEL GROUPS ASK FOR FIRST

1	**Budget controls** A booking agent loops through a holiday sale 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 front-desk agent sends a guest's whole booking to an AI.	Argent reads each request first and blanks out card, US passport and Social Security numbers, plus reservation codes and 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 guest data, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An auditor or internal audit can check a package of it on their own computer and see whether anything was altered.
4	**Model control** A pricing 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 guest data. Each brand can keep its own list. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Guest data must never leave the group's own network.	Argent runs entirely inside your own data center, 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]

Florida, Fla. Stat. 501.171. Affected guests must be told no later than **30** days after a breach is determined; and the state as well when **500** or more Floridians are affected.[3]

HOW A HOTEL GROUP 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: guest 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 HOTEL GROUP'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 IBM / Ponemon, Cost of a Data Breach Report 2026, July 2026 (hospitality average; shadow AI, all industries). 2 California Privacy Protection Agency, CCPA monetary thresholds, Cal. Civ. Code 1798.155(a), effective Jan 1, 2025. 3 Fla. Stat. 501.171(3)(a) and (4)(a), Florida Legislature, read 2026-11-19.