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

ARGENT MODEL CONTROL

EVERY AGENT PICKS ITS OWN MODEL. APPROVE THEM IN ONE PLACE.

Argent Strategic Intelligence decides which AI models the agents that work through Argent may use. One approved list sits at the gateway: a model nobody approved is refused before anything leaves. Providers retire models on a schedule, and requests to a retired model fail,[1] so an agent that hard-coded a model name simply stops working.

Change models without changing agents. A platform rule runs an approved model in place of the one an agent asked for, even at another provider. A model class lets agents ask by role, such as fast, and moves to the next approved model if one fails. Every decision goes on a signed record.

SEE.
The model each agent asked for, and what ran.

PROVE.
A signed record: asked, ran, rule, estimate.

STOP.
Unapproved models are refused at the gateway.

FIVE CONTROLS IN ONE PLACE

1	**Approved list** An agent calls a premium model nobody reviewed.	No model runs until an admin approves it. The list is layered platform, tenant and agent, and each layer can only narrow the one above. A model that is not approved is refused at the gateway, and the refusal names the layer that said no.
2	**Model Swap** A team keeps asking for a costly model.	A platform administrator writes one rule: when an agent asks for model X, run approved model Y. The agent's code does not change, and the answer comes back in the format it asked for, even across OpenAI, Anthropic, Google Gemini, AWS Bedrock and Azure OpenAI.
3	**Model classes** A vendor retires a model your agents rely on.	Agents ask by role, such as fast. An admin points the class at an ordered list of approved models and changes it in one place; every agent that asks for fast moves with it, with no code change.
4	**Failover** One provider has an outage mid-morning.	When an agent asks through a model class and a model fails before it answers (an error, a rate limit, a timeout, a slow first token), the request moves to the class's next approved model, and every check runs again.
5	**Budget-aware swaps** An agent nears its monthly AI budget.	A rule can wait for a budget: at 80 percent by default, it moves the agent to an approved lower-cost model or class. At 100 percent the budget still stops it. Cost is metered at the price of the model that actually ran.

WHY IT MATTERS NOW

Models retire on a schedule. Anthropic: requests to retired models fail.[1] OpenAI: a shut-down model no longer accepts requests through the standard API.[2]

Access is the gap. **92%** of organizations with an AI-related breach lacked proper AI access controls, and incidents involving unapproved AI reached **43%** of the breached organizations studied.[3]

HOW A TEAM STARTS

1. Watch. Argent records what it would refuse or swap, and refuses or swaps nothing until you enforce.

2. Approve. Build the approved list, the first swap rules and classes, then switch on Enforce.

3. Prove. Each record shows the model asked, the model that ran, the rule or class, and why.

BOOK A DEMO. WE WILL SHOW YOU ARGENT CHOOSING MODELS ON YOUR OWN AI TRAFFIC.

NYEDIS.AI/DEMO

Sources. 1 Anthropic, Model deprecations, platform.claude.com, read 2026-11-19. 2 OpenAI, Deprecations, developers.openai.com, read 2026-11-19. 3 IBM / Ponemon, Cost of a Data Breach Report 2026, July 2026.