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

ARGENT AGENT TOOL ACCESS

AGENTS CALL TOOLS AND OTHER AGENTS. YOU DECIDE WHICH.

Argent Strategic Intelligence sits between your AI agents and the tools and agents they call. Tool calls over MCP (Model Context Protocol) pass the MCP Gate; jobs one agent asks of another over A2A pass the A2A gate. Both refuse by default. The MCP and A2A specifications leave authorization to each implementation.[1] [2]

One signed record for both. Every call, allowed or refused, is recorded with who asked, for what and why, with a hash of the arguments, never raw values. On an enforce server a new or changed tool waits for an admin.

SEE.
Which tools and agent jobs each agent can use.

PROVE.
A signed record of every call and refusal.

STOP.
Tool calls and agent jobs nobody granted.

FIVE THINGS IT DOES FOR YOUR TEAM

1	**Given tools only** An invoice agent may read invoices, not send payments.	Each permission names one MCP server and tool, or a whole server, never across tenants. No permission, no call, and the agent's tool list shows only what it can use.
2	**Jobs between agents** The invoice agent asks the payments agent to schedule a run.	With A2A governance on, each job one agent asks of another needs an enabled rule and the permission it names. The caller's identity comes from the connection, never the message.
3	**New tools wait** A server adds a new tool overnight.	On a server in enforce pin mode, a new tool, or one whose definition changed after approval, stays blocked and off the tool list until an admin decides.
4	**Poisoned tools flagged** A tool description hides an instruction to the model.	Argent hashes every tool definition and scans it for injection phrases and look-alike characters, flagging it for an admin, and refuses dangerous-looking arguments, such as command injection, before they reach the server.
5	**Clear refusals** An agent calls a tool it was never given.	The agent gets a refusal in plain words, such as "No matching MCP scope", and nothing reaches the tool or the other agent. If Argent cannot check a call, it says no.

WHAT SECURITY GUIDANCE ALREADY NAMES

OWASP Top 10 for LLM Applications 2025. Names excessive functionality and permissions as root causes of excessive agency; limit the tools agents may call to the minimum necessary.[3]

CVE-2025-6514, rated 9.6 of 10. The mcp-remote connector was open to OS command injection when connecting to untrusted MCP servers, through a crafted reply.[4]

HOW A TEAM STARTS

1. Register. Add your MCP servers, or pull them in from a registry, and switch on A2A governance where agents call agents.

2. Watch. Argent records every call it would refuse while calls pass unchanged.

3. Enforce. Grant each agent its tools and jobs, switch on refusals, and approve new tools as they appear.

BOOK A DEMO. WE WILL SHOW YOU ARGENT DECIDING WHICH TOOLS AND AGENTS YOUR OWN AGENTS CAN REACH.

NYEDIS.AI/DEMO

Sources. 1 Model Context Protocol specification, version 2025-06-18, Security and Trust and Safety. 2 A2A Protocol specification, section 7.5. 3 OWASP Top 10 for LLM Applications 2025, LLM06:2025. 4 NIST NVD, CVE-2025-6514. Shown: traffic routed through Argent, A2A governance on; refusals only in Enforce mode.