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

ARGENT AGENT PERMISSIONS

ANY AGENT, ANY TOOL. UNTIL YOU DECIDE WHO GETS WHAT.

Argent Strategic Intelligence sits between your AI agents and your MCP servers, the programs that give agents tools such as reading a file or sending a payment. An agent can use a tool only if it was given that tool, and any other call is refused with a clear reason. OWASP names excessive permissions as a root cause of excessive agency.[1]

One console page shows every MCP server and tool your agents can reach. You choose who may use each: nobody, everyone, or only those you name. On an enforce server, a new or changed tool waits for an admin, and every call is on the signed record.



SEE.
Every server and tool, and who can use each.



PROVE.
A signed record of every call and refusal.



STOP.
Calls to tools an agent was not given.

FIVE THINGS IT DOES FOR YOUR TEAM

1	**Given tools only** A mortgage agent may look up accounts, not send payments.	Each permission names one server and tool, or a whole server, and never reaches across tenants. No permission, no call, and the agent's tool list shows only the tools it can use. Grants come directly or through a role, group or identity provider.
2	**One page, every tool** The CISO asks who can call send_payment today.	The MCP Gate page lists every server and tool with its state, who holds it by name, and calls and refusals in the last 24 hours. Set each to OFF, BY POLICY or EVERYONE, then allow or deny named people and agents.
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 approves or denies it.
4	**Registries, held** A team already uses the Official MCP Registry.	Argent connects to the Official MCP Registry, the Docker MCP Catalog, or any registry that speaks the same format. Each imported server arrives switched off until an admin approves it.
5	**Clear refusals** An agent calls a tool it was never given.	The agent gets a refusal with a policy id and trace id, never the rule text, and nothing reaches the tool server. The signed record names the server, tool and reason, with a hash of the arguments, never raw values.

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.[1]

Model Context Protocol specification. Says tools represent arbitrary code execution, and that MCP itself cannot enforce its security principles at the protocol level.[2]

HOW A TEAM STARTS

1. Register. Add MCP servers or import them; each imported server arrives off until an admin approves it.

2. Watch. Switch the MCP Gate to Watch: refusals are recorded, calls pass.

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

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

NYEDIS.AI/DEMO

Sources. 1 OWASP Top 10 for LLM Applications 2025, LLM06:2025 Excessive Agency. 2 Model Context Protocol specification, version 2025-06-18, Security and Trust and Safety. Shown: MCP traffic routed through Argent; refusals only in Enforce mode.