One URL.
Every tool call, governed.

Waxell MCP Gateway is the governance layer for AI agent tool use, identity-resolved, policy-checked, fingerprinted, and audited before any upstream sees the call.

THE PROBLEM

High-stakes tool use. Governance still missing.

- ✗ No list of which MCP servers your team has connected
- ✗ Tool calls attributed to a bot, not the person who triggered them
- ✗ No preventive control, governance is after the fact, if it exists at all
- ✗ Offboarding means chasing credentials across 5-15 SaaS consoles

```
{
  "mcpServers": {
    "waxell": {
      "url": "https://mcp-gateway.
<tenant>.waxell.dev/mcp/"
    }
  }
}

// one line replaces every
// upstream MCP config
```

HOW IT WORKS: ONE ENDPOINT, THREE ENFORCEMENT LAYERS.

STEP 01

Add one URL to your agent config

Replace every upstream MCP server entry with a single Waxell gateway URL. No API key in the config file. Works with Claude Desktop, Claude Code, Cursor, Copilot Studio, and any MCP-compatible client.

→ STEP 02

Connect your upstreams

Add GitHub, Slack, Notion, Linear, internal APIs from the catalog. Each upstream is immediately subject to policy rules and fingerprinting. Users connect once via OAuth, the gateway handles token refresh.

→ STEP 03

Governance begins immediately

Every tool on every upstream is hashed at discovery. New tools need admin approval. Tool descriptions are scanned for prompt injection. Every call is identity-resolved, policy-checked, and logged from the first request.

WHAT YOU GET

Identity and Attribution

Every tool call resolved to a real user, not a bot account. On-behalf-of OAuth means GitHub and Slack see the person who triggered the call. One deactivation revokes every upstream grant.

Fingerprints and Prompt Injection Defense

Every tool hashed at discovery. Drifted tools can be blocked until an admin re-approves them. Descriptions scanned for prompt injection at discovery, before agents call them.

Policy Gate

Allow, deny, redact, or require approval, applied before and after the upstream call. Rule changes reach the fleet within 30 seconds. Approvals pause the agent without breaking it.

Audit

Every brokered call logged with identity, decision, and rules that fired. Retention set by your plan. Exportable to CSV. The answer to "show me every tool call last quarter."

Govern the tools your agents are already calling.

One URL. One governance layer. Full audit from day one.

1	5	30s	0
URL per tenant	trust states	policy propagation	gap between run & govern

BUILT FOR: Security Teams · Platform Engineering · GRC / Compliance · IT Operations

WORKS WITH: Claude Desktop · Claude Code · Cursor · Copilot Studio · + any MCP client

WAXELL

waxell.ai/products/mcp-gateway