CEQUENCE

DATASHEET

Your Agentic AI Footprint is Bigger Than You Think

Cequence AI Gateway – AI Discovery

You can't govern what you don't know exists. Business teams are adopting AI agents faster than security teams can track them, and every ungoverned agent is another insider threat with credentials, privileges, and network reach that hasn't been reviewed or approved. These agents aren't breaking in; they're logging in, with valid credentials. A WAF sees an allowed request; an API gateway sees an authenticated API call. Neither maintains state across a session, so a chain of legitimate tool calls that strays outside the agent's intended job looks like a series unrelated, authorized events. Ungoverned and shadow agents create blind spots that turn into expanded attack surface, credential sprawl, and data exposure, and an incomplete inventory that slows the response when something goes wrong. Agentic Zero Trust, the principle behind the Cequence AI Gateway, contains an agent at the boundary of what it can do instead of trying to secure the model itself, and discovery is where that containment starts. Read the research at cequence.ai/agentic-zero-trust.

Cequence AI Gateway – AI Discovery

The Cequence AI Gateway discovers MCP servers, AI agents, LLMs and the users running them in your enterprise, whether they went through an official process or not. It works by analyzing your SIEM logs, so there is nothing to install: no endpoint agent, no MDM rollout, no browser extension, and no inline proxy. Analyzing existing logs means there's no wait for new traffic to accumulate to generate a report.

What One Enterprise Found in Its First Report

A global financial services organization with over 2,000 employees used the Cequence AI Gateway for agentic AI discovery. They found that every employee was using at least one agent. What the security team expected to find, versus what was actually in use, was startling:

	Expected Findings	Actual Findings
Shadow MCP servers	Microsoft 365, Atlassian, GitLab, Docusign, HubSpot	+740%
AI agents	Microsoft Copilot, Claude	+800%
LLMs	Anthropic, Azure OpenAI	+600%

The team's estimate was not careless — those were the reasonable answers, but each was dramatically low.

Cequence AI Discovery at a Glance

- **Maps your full agentic AI footprint:** shadow MCP servers, agents, and LLMs
- **Integrates with your SIEM:** no agents, extensions, or proxies to deploy
- **Immediate first report** with an analysis covering the period you choose
- **Refreshes as often as every four hours** with 30 days of trend history
- **Feeds directly** into the Cequence AI Gateway registries and Agent Personas

What it costs

IBM's Cost of a Data Breach 2026 found that AI-enabled breaches average $6 million and 43% of security incidents involve shadow AI — roughly double the prior year. Of the organizations breached through an AI application, 92% had failed to control access to it, and the most common entry point was a compromised API or plug-in.

How AI Gateway's AI Discovery Works

Connect your SIEM. AI Gateway analyzes existing SIEM logs for AI activity.

Nothing to deploy. No endpoint agent, no browser extension, no traffic re-routing. Read-only SIEM access is all that's required.

Your logs stay where they are. Sensitive data never leaves your SIEM. Only the resulting reports are stored in the Cequence AI Gateway.

Continuous, not a snapshot. Reports refresh on a schedule of your choosing and administrators can trigger an immediate run at any time. See how your AI footprint changes over time.

From Inventory to Control

AI Gateway surfaces what exists and who is using it, and adding those discovered agents, providers, and MCP servers to the Cequence AI Gateway brings it all under governance. AI Gateway surfaces the first four rows below through its discovery. Onboarding those assets into the Cequence AI Gateway unlocks the rest.

	Discovered	Governed by AI Gateway
MCP servers in use	●	●
AI agents in use	●	●
LLMs in use	●	●
Sessions per user	●	●
Individual MCP tool calls	—	●
Tool and schema drift	—	●
Downstream APIs and data reached	—	●
Sensitive data in prompts and responses	—	●
Behavioral containment and enforcement	—	●

Discovery detail varies with the depth of the logs your SIEM collects.

- **Onboard it, then scope it.** Once discovered, MCP servers can be listed the enterprise MCP Registry and appropriately governed, with credentials brokered by the platform rather than held by the agent, and each agent is bound to an Agent Persona — a plain-language job description compiled into actionable policy.
- **Contain it, then prove it.** Agentic Zero Trust judges an agent on what it does across its full session and constrains it the moment it strays. Audit trails in OTEL format export back to the SIEM the discovery data came from.

Governance and Audit Readiness

- **NIST AI Risk Management Framework Govern 1.6 and ISO/IEC 42001 Annex A.4.2.** Both require a maintained inventory of the AI systems and resources in scope.
- **OWASP agentic and MCP Top 10s.** Addresses ASI03 Identity and Privilege Abuse, ASI10 Rogue Agents, MCP09 Shadow MCP Servers, and MCP08 Lack of Audit and Telemetry.
- **EU AI Act.** Article 50 transparency applies now; the Article 26, 14, and 49 deployer duties were deferred to December 2027 — but the inventory they depend on takes longer to build than that.

Built for the Enterprise

Deploy anywhere. SaaS or on-premises, ISO 27001 certified, supporting PCI-DSS, HIPAA, and SOC 2 programs. The same platform that protects many of the world's leading brands now extends this experience to agentic AI governance.

Connect your SIEM and see your agentic AI footprint. Contact us for a **personalized demo.**

CEQUENCE

5201 Great America Pkwy, Suite 240, Santa Clara, CA 95054 | 1-650-437-6338 | info@cequence.ai | www.cequence.ai