CEQUENCE

DATASHEET

Cequence AI Gateway

Security, governance, and control for agentic AI

Cequence built the industry's most effective bot management solution and now secures over 10 billion API calls and 200 million agentic interactions every day. The behavioral intent engine at the core of our platform separates good users, good bots, and fraudulent actors by analyzing intent, not just identity. The AI Gateway extends that proven technology to AI agents: every agent gets a defined job, every action gets verified against that job continuously, and trust is never assumed from only a credential.

Because enforcement lives at the AI Gateway rather than in the model or the endpoint, governance is centralized no matter where agents run: managed devices, cloud platforms, or SaaS agent services like ChatGPT workspaces and Agentforce. Frontier, open-weight, and self-hosted models are all protected. Identity gets the agent in; the AI Gateway governs what it does next and verifies every action against its role.

What Makes the Cequence AI Gateway Different

Agentic Zero Trust Architecture

AI Gateway authenticates agents and then authorizes every action they take. The gateway enforces policy inline with the request path, for the full session, on every tool call. Independent research from Dr. Chase Cunningham and Anthropic arrived at the same architecture Cequence had already built.

Least Privilege Access with Agent Personas

A plain-English job description becomes a default-deny role with per-tool-call permissions: specific MCP tools, API operations, and data objects, nothing else. Personas work immediately from the first tool call, with no learning period required.

Behavioral Intent Monitoring

The Cequence behavioral intent engine analyzes the sequence of actions an agent takes, not simply one call at a time, against per-persona baselines and halts actions that deviate from that baseline.

Sensitive Data Protection

Inline detection, masking, redaction, and blocking on every request and response, with more than 100 built-in detection types supporting PCI-DSS, PHI, SOC 2, and HIPAA compliance.

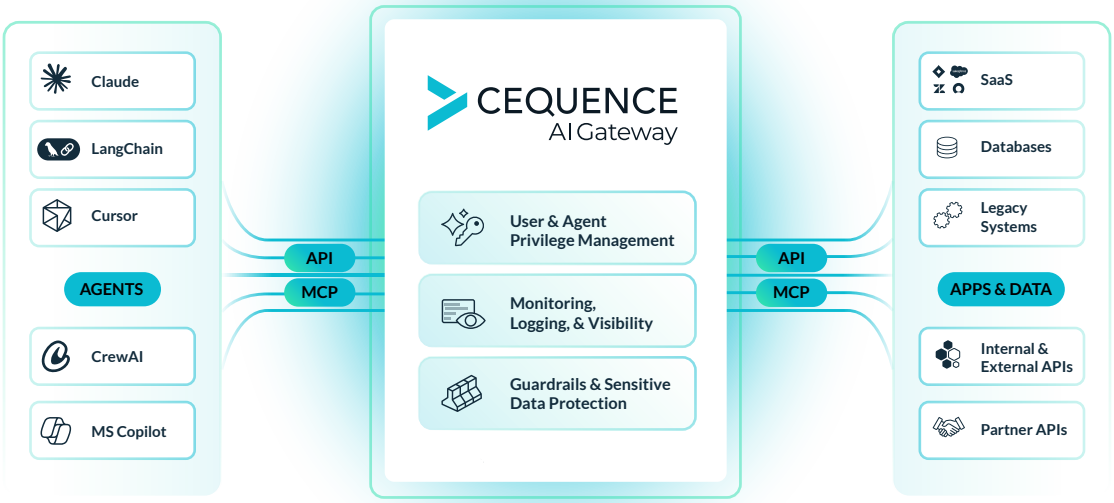

The Cequence AI Gateway provides the security and governance enterprises require to confidently deploy agentic AI workflows at scale.

AI Gateway Capabilities

Agent Identity and Access Management

Agent Personas. Every agent gets a job description that dictates tool access, governed by a central Skill Registry that defines which skills bind to which personas. Each persona maps to a single virtual endpoint and enforcement happens inline at the AI Gateway. Agent Personas expose only the tools the role requires, so the agent receives a short tool list instead of hundreds in every request, which cuts token consumption and improves performance as the model selects the right tool the first time.

Credential Isolation. The agent authenticates to the gateway with a credential that grants access to the gateway and nothing else. The agent and the model never see a backend secret, so a compromised agent or manipulated prompt cannot leak credentials to your systems.

Enterprise Authentication. Multi-IdP support with OAuth 2.1, PKCE, dynamic registration, OIDC integration, and support for legacy authentication, ensuring adherence to organizational identity policies and permissions for both human and non-human identities.

Governance and Compliance

Monitoring and Visibility. Every request records agent, human, tool, time, and outcome, producing the indelible audit trail that compliance frameworks and incident response demand. Data and logs export in OTEL format for SIEM and GRC integration.

Sensitive Data Protection. Actions can pass the permission check and still exfiltratate sensitive data, so the gateway inspects every request and response in flight. It blocks sensitive data moving between allowed tools, flags slow harvesting against per-persona baselines, redacts credentials appearing in tool arguments, and stops production data bound for dev or external tools at egress. Easily integrates with existing DLP infrastructure.

Security

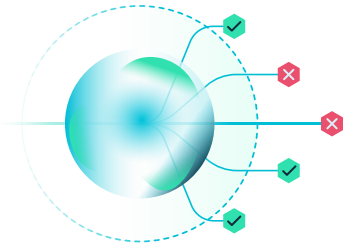

Behavior Anomaly Detection. Behavioral baselines per persona, per user, and per tool catch the patterns single-call inspection misses. An agent reading 200 tickets in a row before opening an email tool would pass every individual permission check; the sequence here is the violation, and the gateway throttles or halts it mid-action based on policy.

Guardrails and Rate Limits. Per-agent, per-tool rate limits with circuit breakers for runaway loops, automated tool risk scoring, and network controls including IP CIDR restrictions, geo-fencing, and IP pinning which requires tokens to be used from the IP address to which they were issued, per persona.

Enablement

No-Code MCP Server Creation. Upload an existing OpenAPI or Swagger spec, or choose from discovered application APIs, and select the endpoints to expose as tools. The gateway generates the MCP server in minutes with no coding, and every server it creates inherits persona scoping, behavioral monitoring, and guardrails. Deploy fully managed in the Cequence Cloud or self-managed with a Helm chart.

Enterprise MCP Registry. Eliminate shadow and rogue MCP servers with a trusted catalog of vetted servers built from official application APIs, plus custom MCPs for your own applications, all centrally provisioned with IAM capabilities. The gateway handles MCP protocol revisions, so no code changes are required as the standard evolves.

Proven in Production

A major global telecom found that a legitimate developer's coding agent ran a task over the weekend, encountered dependencies that prevented it from completing the task, and then attempted 2.5 million tool calls trying to circumvent the roadblock: fabricating file paths, manipulating file SHAs, and probing for write access. Every credential was valid the entire time. The Cequence AI Gateway detected the abnormal behavior, alerted security teams, and generated the complete audit trail and root-cause report for the investigation.

Built for the Enterprise

Deploy as full SaaS with a dedicated tenant per customer, or on premises where your sensitive data never reaches the Cequence control plane. RBAC, continuous environment monitoring, and discrete pre-prod and prod modes are standard, and the platform is ISO 27001 certified, PCI DSS compliant, and maintains a SOC 2 Type II attestation report. Integration with Cequence API Security and Bot Management adds enhanced API specs that boost agent accuracy plus protection from agent-fueled attacks, abuse, and fraud.

Summary

Every enterprise AI roadmap points the same direction: more agents, more autonomy, and access to more applications and data. The Cequence AI Gateway makes that trajectory governable: provision each agent like a new hire with a default-deny role, watch its behavior against that role in real time, and stop the agents whose actions stray from the job description. Cequence co-authored the CIS Model Context Protocol Companion Guide and co-chairs TM Forum's AI-Native Blueprint Initiative, helping define the standards for securing agentic interactions. Whether you are evaluating agent platforms or already running agents in production, Cequence delivers the security, governance, and control that organizations require.

CEQUENCE

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