MCP Security Gateway

Active Defense and Governance for Agentic AI Ecosystems

Inline Protection for Agents

The Model Context Protocol (MCP) is rapidly becoming the default way for AI agents and models to connect to enterprise tools, APIs, and data sources.

As MCP adoption grows, organizations are increasingly exposed to new risks created by tool-connected agents, including unauthenticated tool access, over-permissioned workflows, and prompt injection attacks.

PointGuard AI MCP Gateway provides a secure control point between agents, MCP servers, and enterprise tools, enabling centralized authorization, observability, and policy enforcement.

Key Use Cases

- **Secure agent access to tools and MCP servers**
- **Least privilege to block unauthorized actions**
- **Stop direct and indirect prompt injection**
- **Prevent leaks of sensitive data in workflows**
- **Centralize tool and agent governance**

Discovery and Agentic Inventory

PointGuard AI MCP Gateway provides centralized discovery and inventory of MCP servers, tools, and agents, including how they connect and what capabilities they expose. This lets you:

- **Discover agents, MCP servers, and AI tools**
- **Inventory of tool capabilities and connections**
- **Use a federated registry for share deployments**
- **Enable centralized visibility and governance**

Guardrail Enforcement

PointGuard AI Guardrails are enforced within the MCP Gateway to inspect agent requests and tool calls for prompt injection, jailbreaks, and malicious instructions. Guardrails can be applied at multiple stages of the agent workflow, including pre- and post-filter positions, based on runtime configuration and the tool or data source involved.

- **Detect prompt injection and jailbreak attempts**
- **Support pre- and post-filter guardrail placement**
- **Apply enforcement at multiple points in an agent execution path**

Intent-Based Access Control

A major security gap in MCP environments is that agents can often invoke tools with little or no permission boundaries. PointGuard AI MCP Gateway enables intent based access control at the tool-call level.

This enforces least privilege, including separation of read versus write operations. For example, an agent may be allowed to retrieve data from a CRM tool but blocked from performing update actions unless explicitly authorized.

Access control policies can be applied per agent, per tool, and per operation, enabling granular control over agent behavior and business impact.

- **Enforce least privilege for agents and tool-calls**
- **Separate read, write, and update actions**
- **Apply access rules per agent, tools, and operations**
- **Reduce risk from over-permissioned workflows**

Advanced Threat Protection

The PointGuard AI Platform protects against agent threats that arise when models interact with enterprise tools and data. It blocks stored prompt injection hidden in enterprise content before it reaches the model and reduces risk from unauthenticated localhost MCP servers that can be discovered and exploited by browser-based attacks.

- **Block stored prompt injection in content**
- **Scan tool-retrieved content**
- **Secure local MCP servers**
- **Stop browser-based exploitation**

www.pointguardai.com info@pointguardai.com

MCP Security Gateway

AI Data Loss Protection

PointGuard AI MCP Gateway includes DLP enforcement to prevent sensitive data from leaving the enterprise through AI agent workflows. Even if an injection attempt succeeds, outbound DLP policies can block exfiltration, providing a second layer of defense beyond prompt filtering

- **Stop sensitive data leaks though agent and tools**
- **Block prompt injection exfiltration attempts**
- **Reduce risk of data leaks through agent actions**
- **Enforce consistent DLP policies across workflows**

Security-Aware Observability

The MCP Gateway provides observability into agent-to-tool, and agent-to-agent traffic, through OpenTelemetry-style traces without requiring customers to deploy SDKs inside each agent.

- **Generate OpenTelemetry traces at the gateway**
- **Complete agent-to-tool and agent-to-agent visibility**
- **Support troubleshooting, auditing, and anomaly detection workflows**

REST API to MCP Tool Virtualization

The solution helps scale tool adoption with virtualization of existing enterprise REST APIs into MCP-accessible tools. This allows teams to expose API-based systems to agents through MCP without rewriting each integration.

- **Turn existing REST APIs into MCP-enabled tools**
- **Connect tools without rewriting legacy APIs**
- **Enforce policies across new and existing services**
- **Scale MCP adoption across large environments**

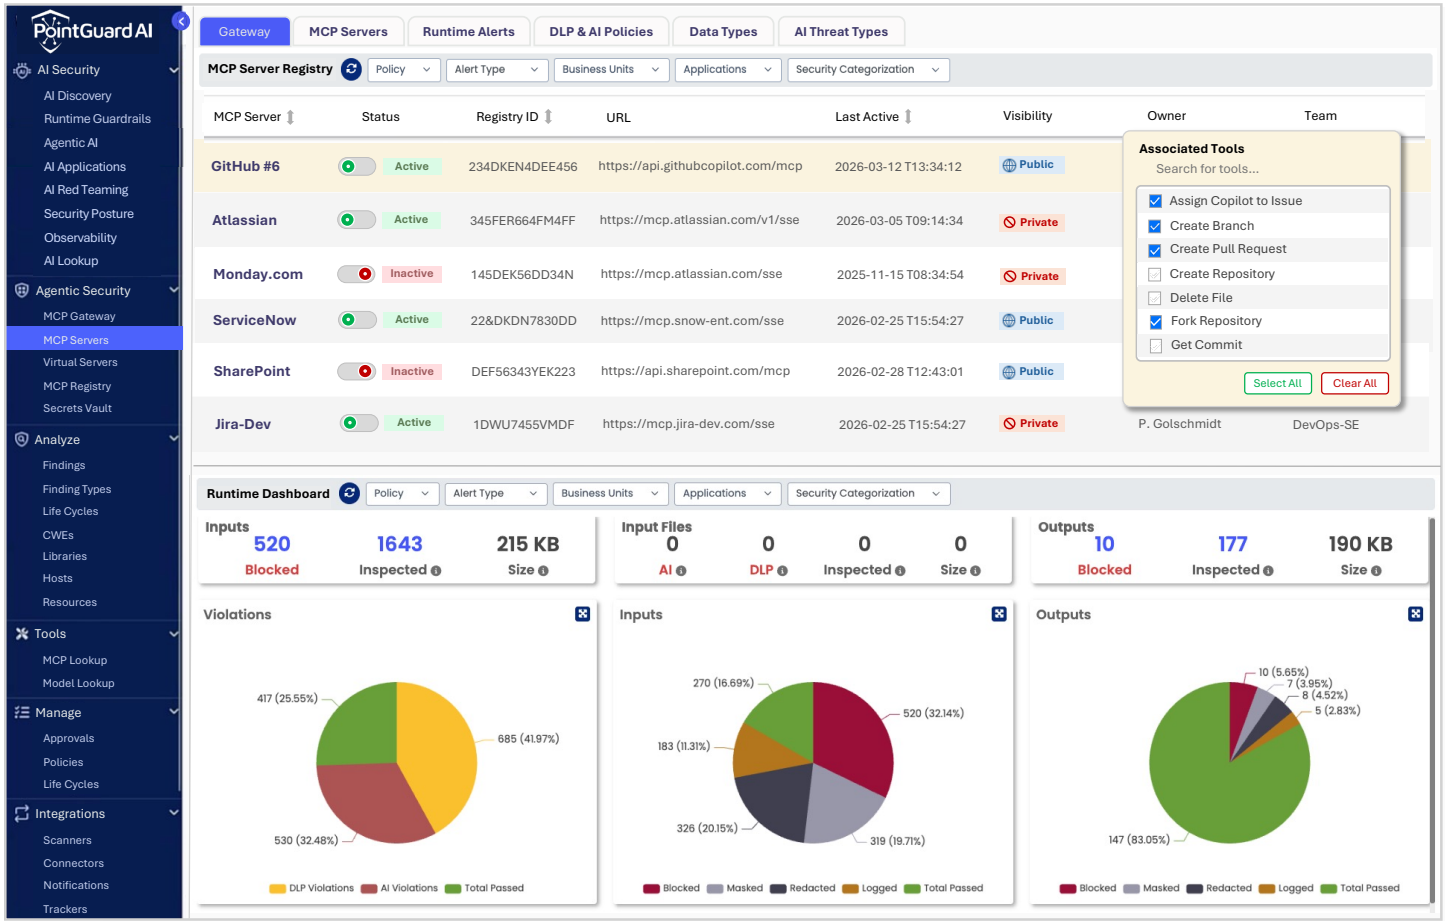

MCP Gateway and Guardrail Dashboard

Hybrid Deployment Options for Data Sovereignty

PointGuard AI adapts to diverse security architectures, regulatory environments, and performance needs. Whether organizations require rapid deployment or strict data residency controls, PointGuard AI provides flexible models that preserve sovereignty, maintain control, and optimize performance for AI agents and tool interactions.

- **SaaS Platform provides fastest time to value**
- **Hybrid data plane in customer cloud ensure sovereignty**
- **MCP Gateway in customer cloud provides complete control**
- **Localized traffic reduces latency and protects sensitive data**

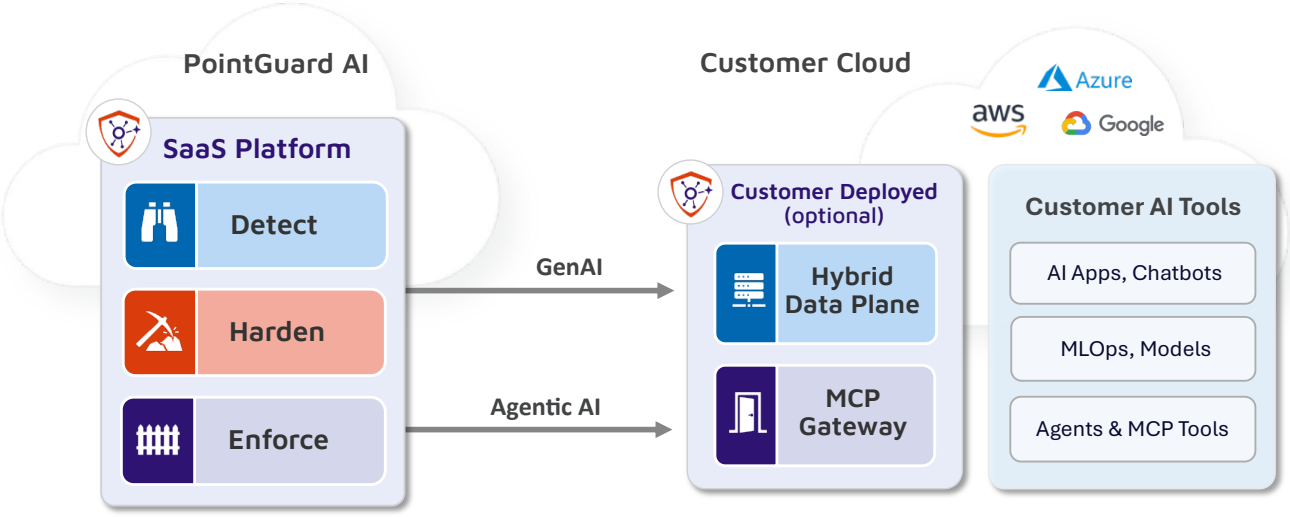

Part of the PointGuard AI Platform

PointGuard AI uniquely secures both AI systems and software applications. Through a single, unified management console all components work seamlessly together to secure the complete AI lifecycle, from discovery to data protection.

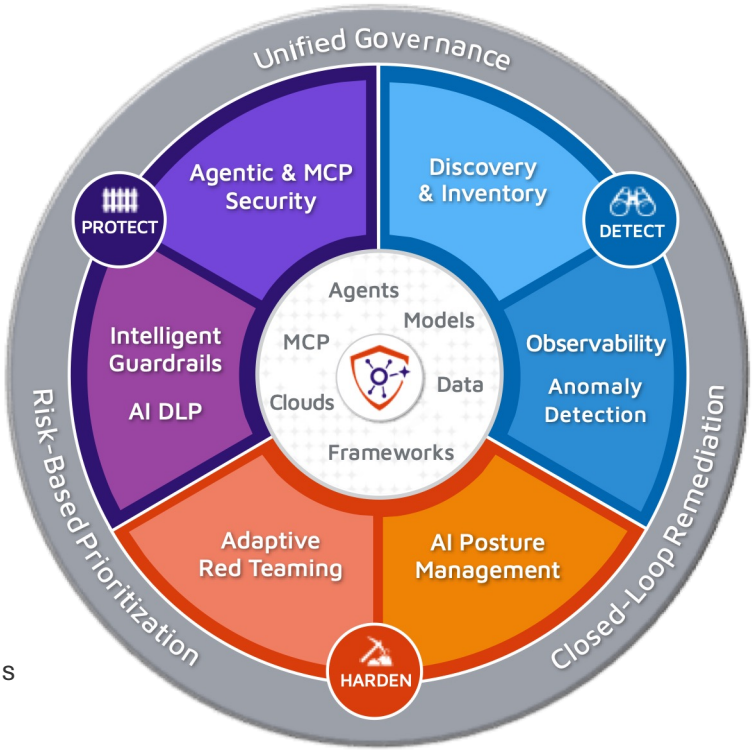

Built for the Agentic Era

As AI systems evolve into autonomous, interconnected agents, security testing must evolve with them. The PointGuard AI MCP Gateway deliver context-aware, policy-driven protection designed for today's AI systems and tomorrow's agentic environments.

Get Started

View demos and detailed technical content on our website or schedule a call to discuss your specific needs with our security experts.

www.pointguardai.com/contact