# MCP Security Gateway

Access Control, Guardrails, DLP, and Governance for Agentic AI

## The Agentic Security Challenge

The rapid adoption of MCP enables AI agents to autonomously access systems, APIs, and data across the enterprise, unlocking powerful new efficiencies and automation. But agent interactions are dynamic, multi-step, and often executed without clear identity, delegation, or authorization controls. This creates a new class of risk where actions cannot be reliably attributed, controlled, or constrained, and where trust boundaries become increasingly difficult to enforce.

PointGuard AI MCP Security Gateway provides an identity-first, zero-trust enforcement point for agent ecosystems, combining identity, authorization, runtime protection, and data security into a centralized control layer. This ensures agents operate securely, transparently, and in alignment with enterprise policies across complex, multi-agent workflows.

## Key Use Cases

- **Identify agents and manage delegated authority**
- **Enforce least-privilege access control for MCP**
- **Prevent threats that can derail agents**
- **Protect sensitive data across agent workflows**
- **Centralize agent, tool, and policy governance**

## Agentic Security Starts with Identity

AI agents introduce a new identity paradigm. Unlike traditional applications, agents operate across multiple systems and often act on behalf of users, services, or other agents. Without clear oversight of identity propagation and delegation, organizations lose visibility and control over which entities are performing authorized actions.

The PointGuard MCP Security Gateway establishes a unified identity plane for agents, ensuring every action is tied to a verified source and explicit authority. Agent identities are uniquely defined and seamlessly integrated with enterprise IAM systems enabling security teams to:

- **Propagate end-user identity via OAuth tokens**
- **Enforce “on-behalf-of” (OBO) delegation**
- **Track agent and user identity per request**
- **Isolate credentials and shared secrets**
- **Audit delegation chains across workflows**

## Intent-Based Access Control

A major security gap in MCP environments is that agents can invoke tools with inadequate permission boundaries. The MCP Security 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, so you can:

- **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**



www.pointguardai.com info@pointguardai.com

# MCP Security Gateway

## Discovery and Agentic Inventory

The PointGuard platform 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**
- **Access an extensive KB of MCP server risks**
- **Enable centralized visibility and governance**

## Security-Aware Observability

The MCP Gateway provides real-time visibility into agent-to-tool, and agent-to-agent traffic, via OpenTelemetry traces, without requiring customers to deploy SDKs inside each agent. Your team can easily:

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

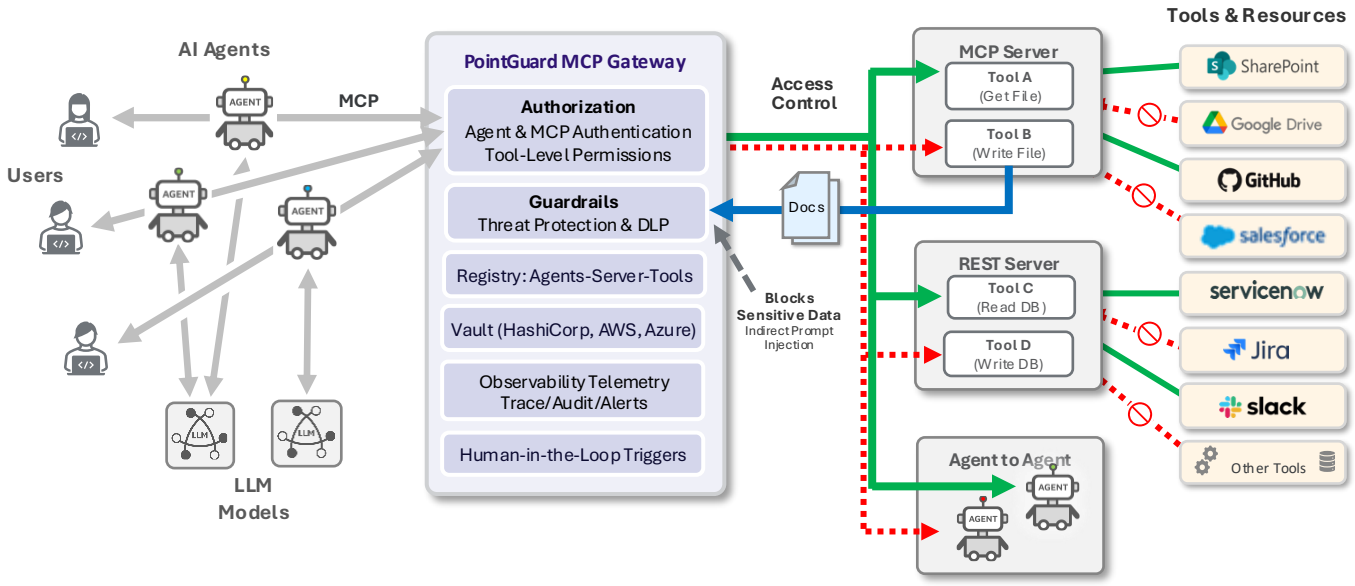


## Intelligent Guardrails

The gateway enforces guardrails to monitor prompts, responses, and tool calls for prompt injection, jailbreaks, data poisoning, and malicious content. Powered by fine-tuned models, controls are applied across workflow stages based on runtime context, tools, and data sources.

- **Detect prompt injection and jailbreak attempts**
- **Filter content pre- and post-execution**
- **Enforce controls across agent execution paths**

## Advanced Threat Protection

PointGuard AI protects against the broader agent attack surface, including prompt injection, jailbreaks, and malicious content designed to manipulate agent behavior. It applies contextual policy enforcement across agent workflows to detect and stop threats in real time.

- **Detect prompt injection and adversarial inputs**
- **Monitor prompts, responses, and MCP activity**
- **Stop exploitation of execution environments**

## Indirect Prompt Injection Protection

PointGuard inspects and controls all retrieved inputs before they reach the model or influence execution, preventing indirect prompt injection from external content, files, code, and metadata that could manipulate agent behavior.

- **Scan retrieved content before LLM ingestion**
- **Inspect tool outputs prior to processing**
- **Block embedded malicious instructions**
- **Prevent unauthorized downstream actions**

## AI Data Loss Prevention

The platform enforces controls to prevent data leaks across agent workflows. Contextual, business-aware controls and automated actions, complying with PCI, PHI, GDPR, HIPAA, and more, to:

- **Prevent sensitive data leaks via agents**
- **Block prompt injection-driven exfiltration**
- **Automate actions (block, mask, redact, notify)**
- **Enforce consistent policies across workflows**

## Human-in-the-Loop Triggers

Some agent actions require explicit human approval. PointGuard enables policy-driven intervention to ensure sensitive or high-risk operations are reviewed and authorized before execution.

- **Trigger policy-based execution pauses**
- **Enable approval workflows (Slack, email, etc.)**
- **Require step-up authentication as needed**
- **Define escalation and timeout policies**

## REST API to MCP Tool Virtualization

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 new policies across existing services**
- **Scale MCP adoption across large environments**

## Secrets Management Vault

Agents require credential access, but poor handling introduces risk. PointGuard integrates with enterprise vaults to enforce identity-based access and eliminate exposure across agent workflows.

- **Enforce identity-based access to secrets**
- **Eliminate static and shared credentials**
- **Prevent secrets in prompts and logs**
- **Support rotation and lifecycle management**

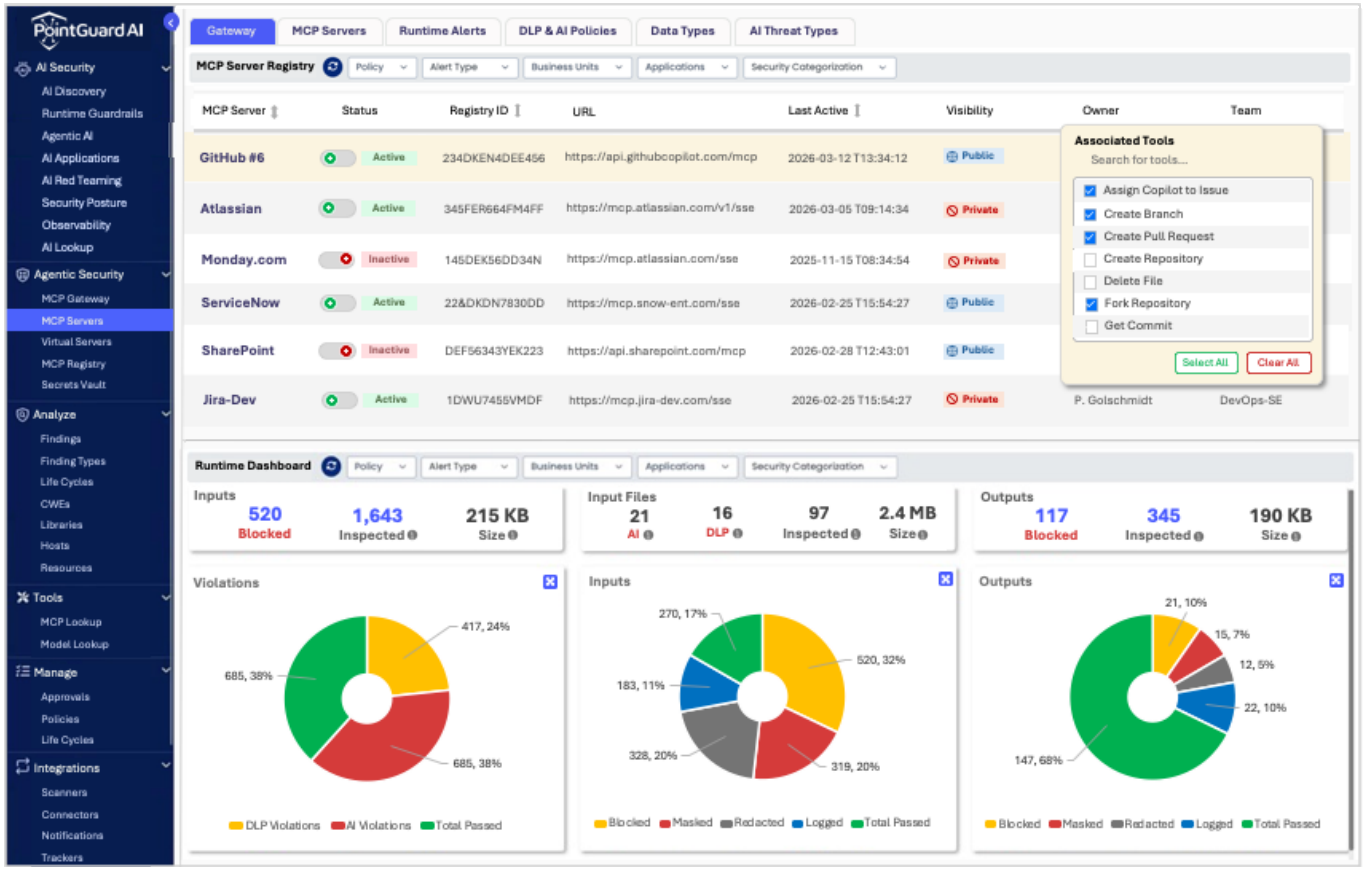


*MCP Gateway and guardrail dashboard*

## Hybrid Deployment Options for Data Sovereignty

The platform 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. This includes:

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

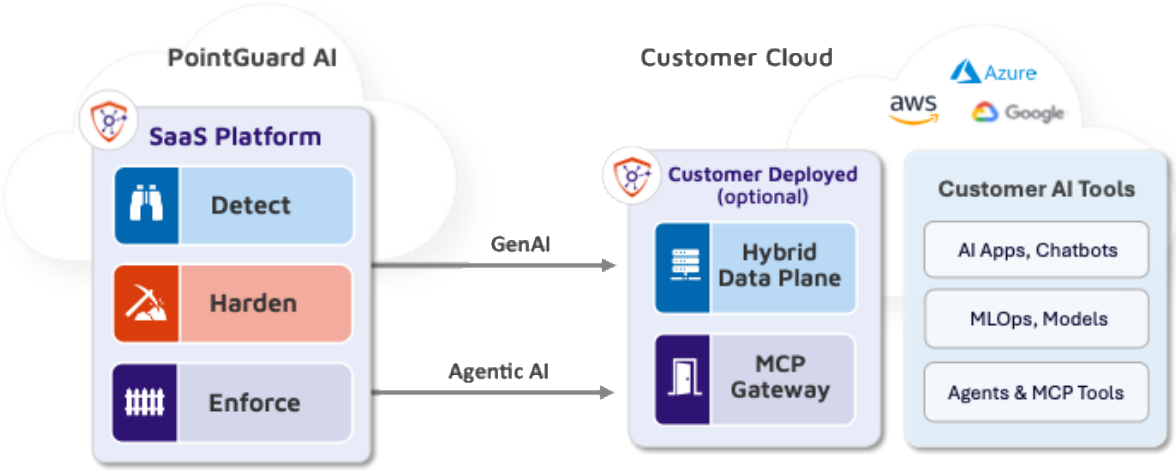


*Flexible platform deployment options*

## Part of the PointGuard AI Platform

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

## Solutions Built for the Agentic Era

As AI systems evolve into autonomous, interconnected agents, security must evolve with them. The PointGuard AI MCP Gateway delivers 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.

Visit: [www.pointguardai.com/demo](http://www.pointguardai.com/demo)

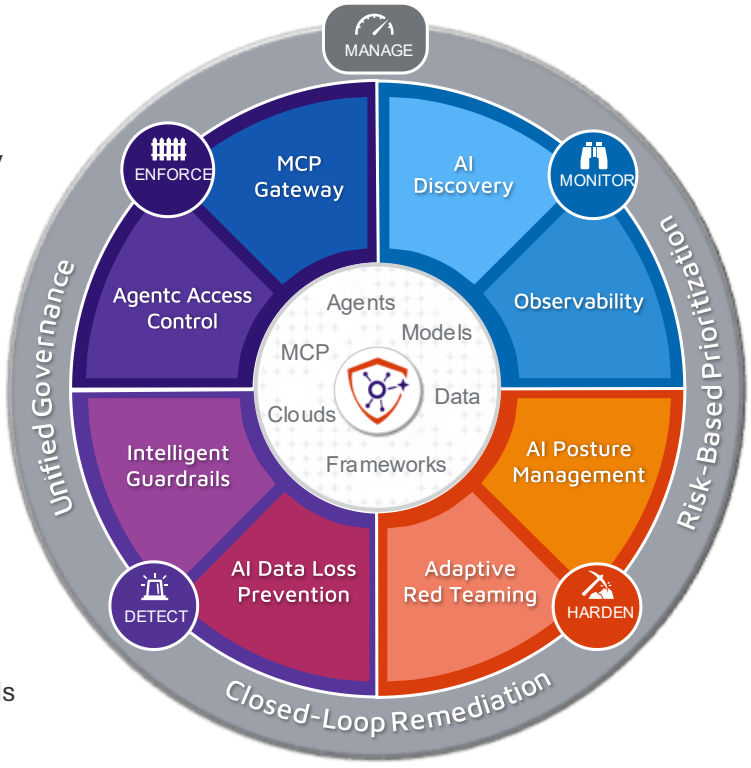


*The full PointGuard AI Agentic Security Platform*

# MCP Security Gateway Capabilities

PointGuard AI

| Agentic Identity & Access Control | |
|---|---|
| **Agent Identity** | Manages explicit dual identities (agent + users) |
| **Authentication** | Integrates seamlessly with enterprise IAM systems |
| **Granular Access Control** | Applies policies per agent, tool, and operation |
| **Read / Write Separation** | Differentiates retrieval vs. modification actions |
| **Approval-Based Workflows** | Requires human approval for high-risk operations |

| Runtime Enforcement | |
|---|---|
| **Agents & Workflows** | Controls agent behavior across tool interactions and execution flows |
| **MCP Servers** | Governs communication between agents and MCP-enabled services |
| **Tools & APIs** | Enforces policies on tool invocation and API-based actions |
| **Agent-to-Agent Interactions** | Monitors and controls cross-agent communication and coordination |

| Context & Policy Engine | |
|---|---|
| **Intent-Based Access Control** | Evaluates the purpose and risk of each action |
| **Business Context Awareness** | Applies policies based on role, environment, and data sensitivity |
| **Dynamic Policy Enforcement** | Adjusts controls in real time based on conditions |
| **Risk-Adaptive Decisions** | Balances security with operational efficiency |

| AI Threat Guardrails | |
|---|---|
| **Prompt Injection** | Detects and blocks direct prompt injection attempts |
| **Indirect Prompt Injection** | Blocks malicious instructions embedded in content, tools, or data |
| **Jailbreak Detection** | Identifies attempts to bypass AI controls |
| **Multi-Stage Guardrails** | Applies controls pre- and post-tool execution |

| AI DLP / Data Protect | |
|---|---|
| **Inline Data Inspection** | Analyzes inputs and outputs in real time |
| **Sensitive Data Detection** | Identifies PII, PHI, and enterprise-sensitive data |
| **Exfiltration Prevention** | Blocks unauthorized data movement through agents |
| **Policy-Based Data Controls** | Enforces consistent protection across workflows |

| Discovery & Agentic Inventory | |
|---|---|
| **Agent Discovery** | Identifies agents across the AI platforms & environments |
| **MCP Server Registry** | Maintains inventory of MCP servers and endpoints |
| **Tool & Capability Mapping** | Maps tool functions, permissions, and dependencies |
| **Ecosystem Visibility** | Provides centralized governance across agentic systems |

| Observability & Monitoring | |
|---|---|
| **Agent-to-Tool Visibility** | Tracks interactions between agents and tools |
| **Agent-to-Agent Visibility** | Monitors multi-agent workflows |
| **OpenTelemetry Tracing** | Generates trace-level observability without SDK deployment |
| **Audit & Forensics** | Supports investigation, compliance, and anomaly detection |

| MCP Ecosystem Enablement | |
|---|---|
| **REST API Virtualization** | Converts existing APIs into MCP-compatible tools |
| **Legacy System Integration** | Connects enterprise systems without rewriting APIs |
| **Policy Enforcement Across Tools** | Applies consistent controls to all integrations |
| **Scalable MCP Adoption** | Enables safe expansion of agent capabilities |

| Advanced Threat Protection | |
|---|---|
| **Stored Prompt Injection Defense** | Blocks malicious instructions embedded in enterprise data |
| **Tool Content Scanning** | Inspects retrieved data before it reaches models |
| **Local MCP Server Protection** | Secures localhost and exposed MCP endpoints |
| **Browser-Based Attack Mitigation** | Prevents exploitation via client-side discovery and abuse |

| Deployment Options | |
|---|---|
| **SaaS Deployment** | Fast time to value with managed infrastructure |
| **Hybrid Data Plane** | Maintains data sovereignty in customer environments for data & policies |
| **Customer-Hosted Gateway** | Customer-Hosted Gateway |
| **Localized Processing** | Reduces latency and protects sensitive data |