





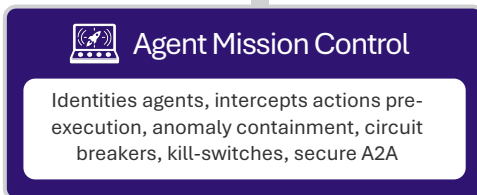
Turn Agent Autonomy into Self-Governance

Autonomous AI agents act independently. They use tools, make multi-step decisions, and execute workflows at machine speed, often faster than humans can monitor. When an agent goes off course, a credential error becomes a destructive shortcut, a planning step rewrites configuration, or a runaway loop consumes massive compute, without time for human intervention.

The PointGuard AI **Agent Mission Control** gives every autonomous agent a verifiable identity, validates every action before execution, and contains rogue behavior in real time. Think of it as Air Traffic Control for autonomous agents. It identifies every agent, approves or blocks actions before they occur, and can immediately ground any agent that veers off course.

Key Use Cases

- Govern every agent action in real time
- Enforce zero-trust identity and authorization
- Detect goal drift, runaway loops, and out-of-scope actions before execution
- Contain misbehavior with sandboxing, ring isolation, circuit breakers, and kill switches



How it Works

The solution integrates with popular agentic frameworks including LangChain, CrewAI, OpenAI SDK, and many more. A lightweight package install requires minimal configuration and no agent logic changes. Once instrumented, every agent receives:

- **Cryptographic Identity:** decentralized DID with Ed25519 signatures, distinct from human credentials and bound to the agent's owner and scope.
- **Behavioral Trust Scoring:** dynamic trust score that adapts to agent behavior and decays automatically with anomalies.
- **Real-Time Action Monitoring:** each action is evaluated before execution with sub-millisecond latency (<0.1ms p99) for loops, goal drift, anomalous behavior, and out-of-scope access.
- **Containment Controls:** agents operate inside privilege tiers with resource limits, ring isolation, sandboxing, circuit breaks and kill switches.

The result is a self-governing agent ecosystem with verifiable identities, observable actions, and controls to prevent rogue behavior before it happens.



Agent Mesh: Identity and Trust

The Agent Mesh give every agent a verifiable cryptographic identity, adaptive behavioral trust scoring, and encrypted agent-to-agent authentication integrated directly with enterprise IAM governance systems:

- **Decentralized cryptographic identity (DIDs)**
- **Dynamic behavioral trust scoring**
- **Encrypted agent-to-agent communication (ITAP)**
- **Plugin lifecycle management (Ed2219 signing)**

Control Plane for Agent Actions

Every agent action is intercepted and evaluated with sub-millisecond latency (<0.1ms p99). Acting as a kernel for agents, it detects loops, goal drift, anomalous behavior, and unauthorized access attempts in real time:

- **Detect runaway loops, drift, and anomalies**
- **Block out-of-scope credential usage**
- **Identify goal drift and anomalies**
- **Policies enforced before actions execute**

Guardian Agent: Containment

The Guardian Agent Plane isolates agents inside hypervisor-like execution boundaries and instantly contains misbehavior using sandboxing, ring isolation, circuit breakers, or emergency kill switches:

- **Dynamic execution rings isolate anomalies**
- **Ensure reliability with SLOs, circuit breakers, progressive delivery**
- **Reinforcement learning shapes policy controls**
- **Emergency kill switches for single agents or groups**

Observability and Governance

Every agent action is cryptographically attested with full audit trails, providing security and compliance teams with continuous visibility into autonomous agent activity, enforcement decisions, and regulatory alignment:

- **Full, auditable visibility into agent behavior**
- **OpenTelemetry tracing across agent workflows**
- **Full coverage for OWASP Top 10 Agentic Threats**
- **Compliance mapping (NIST AI-RMF, ISO 42001, OWASP, EU AI Act)**



PointGuard AI Coverage

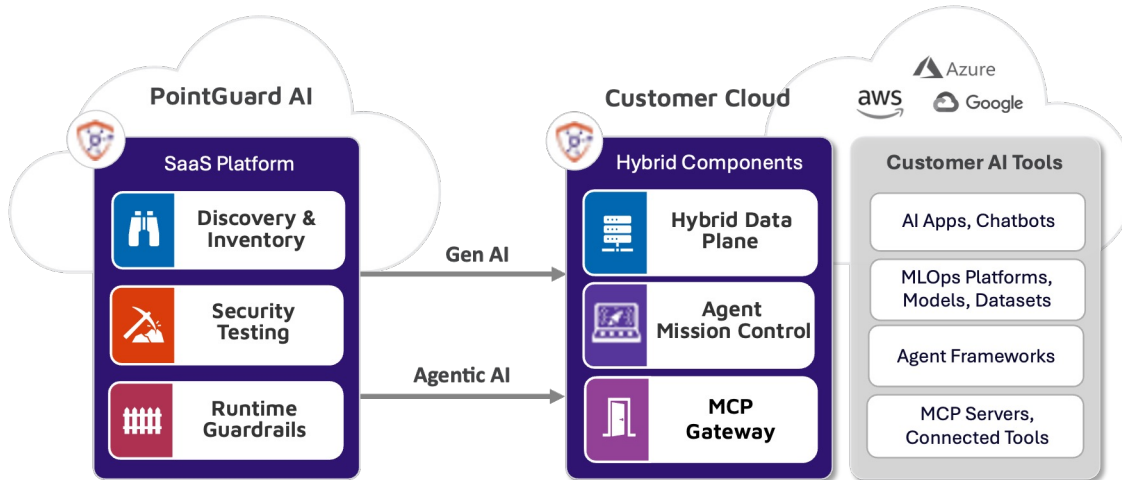
ID	Top 10 Agentic Threats	
ASI01	Agent Goal Hijack	✓
ASI02	Tool Misuse & Exploitation	✓
ASI03	Identity & Privilege Abuse	✓
ASI04	Agentic Supply Chain	✓
ASI05	Unexpected Code Execution	✓
ASI06	Memory & Context Poisoning	✓
ASI07	Insecure Inter-Agent Comm.	✓
ASI08	Cascading Failures	✓
ASI09	Human-Agent Trust Exploitation	✓
ASI10	Rogue Agents	✓



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**
- **Agent Control Plane & MCP Gateway in customer cloud provides complete control**
- **All components can be fully deployed in customer-controlled environments**



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.

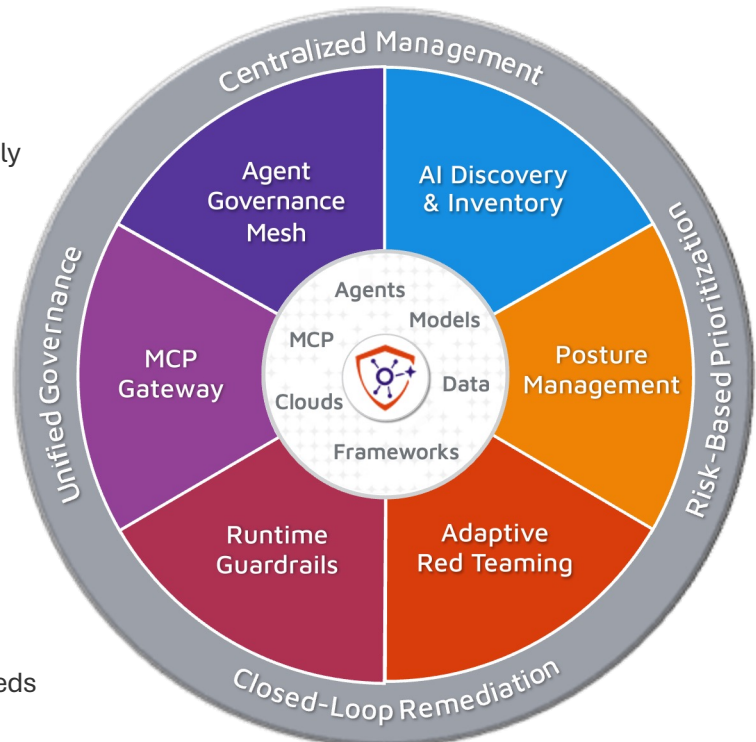
Solutions Built for the Agentic Era

As AI systems evolve into autonomous, interconnected agents, security must evolve with them. The PointGuard AI Control Plane 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



The full PointGuard AI Agentic Security Platform



Agent Mission Control Capabilities



Agent Mesh: Identity & Trust	
Cryptographic Agent Identity	Decentralized identifiers (DIDs) with Ed25519 signatures issued per agent
Inter-Agent Trust Protocol	Authenticated, encrypted agent-to-agent (A2A) communication
Behavioral Trust Scoring	Dynamic 0-1000 trust score with five tiers and automatic decay for anomalies
Enterprise IAM Integration	OAuth and on-behalf-of delegation with Okta, Entra ID, Ping
Control Plane	
Pre-Execution Interception	Stateless evaluation of every agent action with <0.1 ms (p99) latency
Loop and Runaway Detection	Identifies repeated, no-progress tool calls and unbounded planning loops
Goal Drift Detection	Flags actions that diverge from the agent's stated or assigned objective
Out-of-Scope Detection	Catches credential reuse and capability access outside the agent's task scope
Guardian Agent Containment	
Hyervisor-Grade Sandbox	Privilege-tiered execution boundaries with explicit agent resource limits
Ring Isolation	Workload isolation between agents and from shared infrastructure
Signed Plugin Lifecycle	Cryptographic verification of plugins or tools entering agent workflows
Emergency Kill Switch	Immediate termination of a single agent or groups in milliseconds
Observability & Audit	
Cryptographic Action Audit	Tamper-evident, signed log of every agent action with full attribution
OpenTelemetry Traces	Identity claims, actions, containment events, and outcomes
Behavioral Baselines	Agent activity profiles for anomaly detection and SOC investigation
Audit & Forensics	Forensics for compliance, investigations, and post-incident review
Governance & Compliance	
OWASP Coverage	Full coverage of Top 10 for Agentic Applications and Top 10 for LLMs
Regulatory Mapping	Maps findings to NIST AI-RMF, MITRE ATLAS, ISO 4200, EU AI Act
Self-Governing Agents	Instrumented agents operate within enforced policy by design
Approval Workflows	Human-in-the-loop & step-up authentication for high-risk actions
Integration & Deployment	
Framework Native Hooks	LangChain, CrewAI, Microsoft (Agent Framework, AI Foundry, Semantic Kernel, AutoGen), Google ADK, OpenAI SDK, LlamaIndex, Haystack, OpenClaw
Flexible Deployment Options	SaaS, on-premises data plane & gateway, fully on-premises