



Product Datasheet



SaaS

AI Agent Governance & Security

Discover, Govern, and Secure AI Agents Across
Your Enterprise

For enterprise security, IT operations, compliance, and infrastructure leadership teams.

Product Overview

Platform Goal	Deployment Model	Primary Security Outcome
Discover, assess, test, control, and audit AI agents across cloud agent platforms and managed endpoints, and govern access to the MCP servers in use across the organisation.	SaaS.	Bring unknown and ungoverned AI agents under visibility, establish ownership, control agent activity at runtime, and produce audit evidence of agent behaviour.

Securden AI Agent Governance & Security provides a central inventory of AI agents running across cloud agent platforms and managed endpoints. It maps each agent to the models, tools, connectors, and workflows behind it, scores its risk, generates attacks to find weaknesses, and applies policy to its activity at runtime. On endpoints it covers AI applications, browser-based AI, coding assistants, and local LLM runtimes. It also manages the MCP servers in use and controls which users can call which tools.

Core Capabilities

Capability Area	Technical Functionality
AI Agent Discovery and Mapping	Discovers agents across connected cloud platforms and managed endpoints. Records who owns each agent, and maps each one to the models it calls, tools it can invoke, connectors it uses, and workflows it participates in.
Shadow AI Detection and Control	Detects unsanctioned AI on managed devices through the Securden endpoint agent: installed applications, browser-based AI, coding assistants in any IDE, local LLM runtimes, and traffic to AI services. Policies define which tools are permitted.
Risk Assessment	Calculates a risk score for each agent from its guardrail configuration, the access permissions it holds, its monitored runtime activity, and red teaming results, and ranks agents by risk.
AI Red Teaming	Generates attacks from each agent's own instructions, tool access, and guardrails, rather than executing a fixed test set.

Capability Area	Technical Functionality
MCP Server Management	Inventories the MCP servers in use, the tools each exposes, and their owners.
MCP Access Control	Governs which users can call which tools, and blocks calls outside policy.
Runtime Security	Applies eight security controls to agent activity at runtime, enforced at the gateway. Actions are Block, Monitor, and Redact.
Policy and Governance	Assigns shadow AI, runtime, and MCP policies to agents, agent groups, and endpoints.
Audit	Records agent and MCP activity, and administrative actions.

Platform Support

Source	Source Type	Discovery	Runtime Security
Copilot Studio	Cloud agent platform	Yes	Yes
Bedrock	Cloud agent platform	Yes	Yes
Azure Foundry	Cloud agent platform	Yes	Yes
AgentCore	Cloud agent platform	Yes	Yes
Endpoint AI agents	Endpoint	Yes	Yes

AI Agent Inventory

Object	Information Captured
Agent	Platform, owner, agent instructions, status, last activity, risk score
Models	Models called by the agent, such as Claude, GPT, or Amazon Nova

Object	Information Captured
Tools	Tools the agent can invoke, such as a search tool, an email tool, or a database query tool
Connectors	Applications, systems, and data sources the agent is connected to, such as SharePoint, Salesforce, or an internal database
Flows	Automated workflows the agent runs in, such as an agent triggered whenever a support ticket is created
Findings	Security issues associated with the agent
Policies	Runtime security policies applied to the agent
Activity	Recorded agent activity

Discovery and Risk Scoring

Discovery and Risk Scoring	Detail
Cloud platform discovery	API connection to the customer's platform account. No installation required.
Endpoint discovery	Securden endpoint agent installed on managed devices, deployed manually or through Group Policy (GPO).
Discovery execution	On demand or on schedule
Risk analysis	Risk scores are calculated by Securden from the agent's guardrail configuration, access permissions, monitored runtime activity, and red teaming results.

Runtime Security Controls

Control	Detects	Actions
Prompt Injection	Prompt injection attacks	Block, Monitor
Personal Information	Personal and sensitive information	Block, Monitor, Redact
Secrets	Exposure of secrets and credentials	Block, Monitor, Redact
Toxicity	Toxic content	Block, Monitor
Known Attack Pattern	Known attack patterns	Block, Monitor
Dangerous Command	Potentially dangerous commands	Block, Monitor
Harmful Content	Harmful or unsafe content	Block, Monitor
Malware	Malware and malicious content	Block, Monitor

Area	Detail
Enforcement	All controls are enforced at the gateway. The Securden endpoint agent detects AI activity on managed devices and routes it to the gateway.

AI Red Teaming

Area	Detail
Test method	Attacks generated from the agent's own instructions, tool access, and guardrail configuration
Execution	On demand and on schedule

MCP Servers

Model Context Protocol (MCP) servers connect AI tools and agents to enterprise systems. Each server exposes a set of tools, and each tool is a specific action such as creating a repository, editing an issue, or reading a page.

MCP Server Inventory	Detail
Server inventory	MCP servers in use across the organisation, each with its owner, tool count, and status
Adding servers	Added from a catalogue of supported providers
Tool discovery	Tools available on each server are listed automatically once the server is added
Supported providers	GitHub, Atlassian, Notion, GitLab, AWS IAM, ClickUp, Postman, Asana, Slack, Semgrep, Linear, CircleCI, Postgres, MongoDB, Zoom, and Sentry

MCP Control Policies	Detail
What a policy controls	Tool access on a given MCP server, by user
Scope	All users, or named users
Enforcement	At the gateway. Tool calls outside policy are blocked
Monitoring	Tool calls over time, calls allowed and blocked, most used servers, most blocked tools, and server health

Policy, Audit, and Reporting

Area	Detail
Policy types	Shadow AI policy for managed endpoints; Runtime Security policy for AI agents on connected platforms; MCP control policy for tool access

Area	Detail
Assignment	Individual agents, agent groups, endpoints, and users
Audit trails	Every agent and MCP operation, and administrative actions in the platform. Each entry records the user, source IP, and time.
Prebuilt Reports	Agents lacking guardrails, hardcoded credentials, AI dependencies, Copilot Studio agents and owners, red teaming, critical activities, and blocked commands. MCP: tool interactions by user, servers and tools accessible per user, server inventory, and server health.
Export formats	PDF, CSV and XLSX
Audit retention	Indefinite

Architecture and Deployment

Architecture and Deployment	Detail
Deployment model	SaaS
Hosting	AWS
Agent requirement	Agentless for cloud platform discovery. The Securden endpoint agent is required to detect AI activity on managed devices and route it to the gateway.
Cloud platform permissions	Permissions to integrate the gateway and to read AI agent logs. Specific permissions vary by platform.
Gateway	Hosted by Securden. No customer-side installation required
Required outbound connectivity from endpoints	HTTPS 443 outbound. No inbound firewall rules required

Endpoint Agent

Area	Detail
Function	Discovers AI applications, browser-based AI, AI coding assistants, and other AI activity on managed devices, and routes it to the gateway for policy enforcement
Supported operating systems	Windows 10 and later; Windows Server 2019 and later
Deployment methods	Group Policy, or manual installation
Proxy support	Supported

Security and Encryption

Area	Detail
Encryption at rest	AES-256
Encryption in transit	TLS 1.2 or later
Key management	Per-tenant keys managed through AWS KMS
Tenant isolation	Customer data segregated per tenant with a unique encryption key
Data residency	Region selected at provisioning: United States, European Union, or India
Certifications	ISO/IEC 27001 and SOC 2 Type II

Identity, Access, and Integrations

Identity, Access, and Integrations	Supported Options
Platform authentication	Local accounts, Active Directory, Microsoft Entra ID, LDAP, and RADIUS
Single sign-on	SAML 2.0 for Okta, Microsoft Entra ID, ADFS, OneLogin, PingIdentity, Google Workspace, and other SAML identity providers
Directory-based authentication	Authenticates through the configured directory; directory user passwords are not stored
User lifecycle	Directory integrations support onboarding, offboarding, and access provisioning
Multi-factor authentication	TOTP apps including Google Authenticator and Microsoft Authenticator, Securden Authenticator, Duo Security, RSA SecurID, YubiKey, email OTP, and RADIUS-based MFA
Role-based access control	Predefined and custom roles
SIEM integration	Supported
Supported browsers	Chrome, Edge, Firefox, and Safari

Commercial Notes

Area	Detail
Licensing	Annual subscription, priced per discovered AI agent. Endpoint agent licences purchased separately
Free trial	14 days, no credit card required
Non-production environments	No separate licence required
Support	24x5 standard; 24x7 premium

Common Use Cases

Use Case	Description
AI Agent Inventory	Discover agents across cloud platforms and managed endpoints, and maintain a central record of each one with its platform, owner, instructions, activity status, models, tools, connectors, and workflows.
Shadow AI Discovery and Control	Identify AI applications, browser-based AI, coding assistants, and local models in use on managed devices, and enforce which are permitted.
Agent Ownership	Identify who is accountable for each agent and surface those with no owner, so they can be assigned, reviewed, or removed.
Agent Access Mapping	Establish what each agent can reach through its models, tools, connectors, and workflows, and identify those with broader access than expected.
Risk-Based Remediation	Rank agents by risk score so investigation and remediation start with the most exposed rather than the most recently noticed.
Agent Security Testing	Generate attacks against individual agents to find weaknesses in their instructions, tool access, and guardrails. Re-run as an agent changes.
Runtime Attack and Data Protection	Inspect activity for prompt injection, known attack patterns, malware, dangerous commands, personal information, and credentials, and block, monitor, or redact before the request reaches the model.

Use Case	Description
MCP Tool Access Control	Restrict which users can call which tools on each MCP server, so a user with access to a system is not automatically granted every action that system allows.
MCP Access Review	Report on which users can reach which MCP servers and tools, and evidence it for internal review.
Policy Enforcement at Scale	Group agents by application, business function, environment, or ownership and assign policies at the group level.

See how Securden can help your organization

[Book a Demo](#)



 +1 844 228 0327  +44 203 868 6147

 support@securden.com  www.securden.com