VisionX Governance
BY VISIONX SOFTWARE

VISIONX AI SHIELD

Put controls between AI agents and the tools they use.

AI Shield is an **enforcement-ready gateway you deploy in your own network**. It discovers agents, checks every MCP tool call against your policies, holds risky actions for approval and keeps a signed, tamper-evident audit trail. Supports MCP 2025-11-25.

governance.visionx.software

Request a demo: visionx.software/contact

VisionX Software, Inc. · 16192 Coastal Highway, Lewes, Delaware 19958, USA

THE PROBLEM

Agents now act, often with no controls

- AI agents call MCP tools that read records, send messages and change systems, usually with a broad service credential and no human in the loop.
- Security teams rarely know which agents exist, which tools they reach, or what arguments they send.
- When something goes wrong, there is no reliable record of who asked, what was called and why it was allowed.

FOUR ENGINES

One Shield, four engines

01

Agent discovery & inventory

Find the AI agents, MCP servers and tools in use, record owners and environments, and flag agents not linked to a governed record.

02

MCP security gateway

An enforcement-ready gateway you deploy in your own network. Agents connect through it to approved MCP servers, with MCP 2025-11-25 support.

03

Tool-call firewall & policy engine

Evaluate each tool call against published policies: allow, deny, redact or hold for human approval, with argument-level rules.

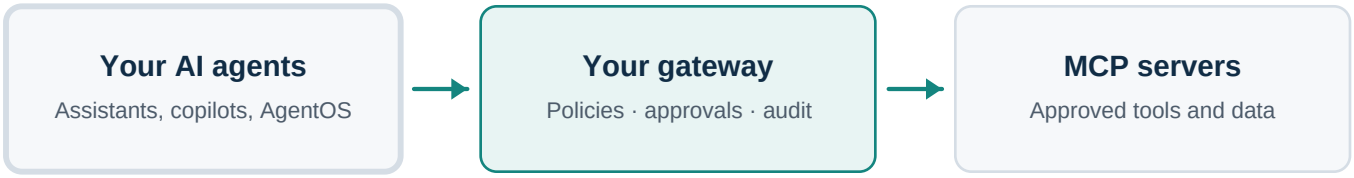
04

Audit, replay & investigations

Every decision lands in a signed, hash-chained audit trail. Replay is display-only, so investigations never re-run a call.

HOW IT WORKS

Agents -> your gateway -> MCP servers



Policies decide allow, deny, redact or hold. **Approvals** issue single-use, argument-bound grants reviewed by a different person. **Audit** records every decision in a signed chain. Start in observe mode, then turn on enforcement when ready.

KEY CAPABILITIES

What security teams get

No-code rules

Build tool and data rules from templates, then simulate before publishing.

Approval center

Single-use grants bound to the exact arguments, actor and tool.

Data-leak detection

Secrets, credentials, personal data and national ID numbers; redacted evidence only.

Prompt-injection & tool-poisoning checks

Inspect tool descriptions and responses for hidden instructions and schema drift.

Signed, tamper-evident audit

Per-tenant hash chain with signed checkpoints you can verify.

SIEM & webhook delivery

Send decisions to your SIEM or signed webhooks with retries.

Gateway setup wizard & health alerts

Copy-ready setup, connectivity checks and alerts when a gateway stops reporting.

AgentOS integration

Register AgentOS agents as metadata records with owner approval.

Detection is layered, not a guarantee.

COMPLIANCE MAPPING

Mapped frameworks and editions

AI Shield controls and evidence are mapped to these frameworks to support your assessments. **These are mappings, not certifications.** SOC 2 audit in progress.

NIST AI RMF | NIST CSF 2.0 | ISO/IEC 27001 | ISO/IEC 42001 | SOC 2 | OWASP Top 10 for LLM & Agentic Apps | Model Context Protocol (MCP)

EDITIONS

Start free, grow into your own gateway

Developer	Business	Enterprise	Private / Regulated
Free	**From $12k/yr**	**From $48k/yr**	**From $120k/yr**
Agent inventory and policy checks for small teams.	Hosted gateway, observe mode, exceptions and core inspection.	Full Shield suite: firewall, signed policies, fleet and red-team harness.	Everything in Enterprise plus the self-hosted gateway in your network.

The self-hosted gateway is available in Private / Regulated only. Annual USD starting prices.

See AI Shield with your own agents

Start free at governance.visionx.software · Request a demo at visionx.software/contact

Take the tour at governance.visionx.software/tour