AVIATRIX

SOLUTION BRIEF

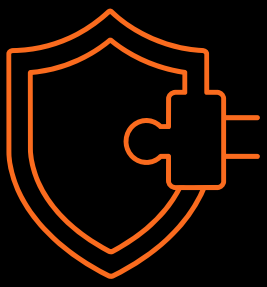

Contain Enterprise AI Agents

Powered by Aviatrix + Microsoft — Agent Control Specification Edition

This Containment Plugin compiles Microsoft Agent Control Specification policy into enforced network rules for AI agents.

The Threat

Enterprises run multiple AI agent frameworks at once – LangChain, AutoGen, CrewAI, Semantic Kernel, Microsoft Copilot Studio, and vendor-shipped copilots – across Kubernetes, serverless, and managed cloud agent services. Every agent reaches Model Context Protocol (MCP) servers, large language model (LLM) APIs, and internal endpoints over the network. Existing controls were never designed to govern that traffic by agent identity.

Agent governance frameworks stack the same problem CI/CD runners do: they define what an agent may reach in code, but cannot guarantee every agent process in the estate loads and honors that policy. A vendor-shipped copilot has no integration point, a legacy workload was never wired in, and a compromised agent can simply ignore its own policy file, but each still reaches the same MCP servers and LLM APIs, unguarded. Constraining that reach at the network layer materially reduces the **Blast Radius** of a compromised or ungoverned agent. Domain-based enforcement is well-suited to this case because it evaluates the destination rather than the caller's identity: blocking an unauthorized tool call regardless of which agent process initiated it.

The Architecture

Aviatrix governs what AI agents can reach. An open-source plugin compiles a Microsoft Agent Control Specification (ACS) policy file directly into Aviatrix Distributed Cloud Firewall (DCF) rules, enforced at the network with zero bypass risk, as enforcement occurs outside the agent trust boundary. No agent code changes are required.

How it works in practice

Coverage	Enforcement Method	Bypass Risk	Status
Any agent on any framework, cloud, or runtime	ACS policy compiled to DCF rules at the Spoke Gateway	Zero (network-layer; no process to kill)	Architecture Validated

Three Things Your Current Stack Can't Do

01 FQDN-based default-deny egress per agent

Compile an ACS policy file into DCF WebGroups that allow only the MCP servers and LLM endpoints an agent is authorized to reach – by domain name, not ephemeral Internet Protocol address – and deny everything else. An agent reaching for an unauthorized endpoint is blocked at the network layer before a single byte leaves your environment, even if the agent never loaded its policy file.

02 Enforcement outside the agent trust boundary

DCF evaluates every outbound connection at the gateway: no sidecar, no mesh dependency, no process on the workload to disable. A vendor-shipped or compromised agent cannot reach an unauthorized tool because the path does not exist. **Blast Radius** shrinks from "any endpoint the network leaves open" to "only the tools this agent was explicitly permitted to use."

03 Enterprise security fabric integration

AI agent policy lives alongside production workload policy, continuous integration and continuous delivery policy, and Model Context Protocol server policy in a single Aviatrix Distributed Cloud Firewall plane. CoPilot provides unified visibility across agents, pipelines, and production workloads from one console, correlating ACS policy events with gateway flow logs by session identifier. No standalone agent governance tool provides this.

Compliance Evidence

For SOC 2, HIPAA, PCI-DSS, FedRAMP, and NIST Secure Software Development Framework environments, post-incident reviews now expect architectural evidence of network egress enforcement for AI agent workloads, not just written policy.

Aviatrix Distributed Cloud Firewall WebGroup and SmartGroup configurations are configured by the the acs-to-dcf policy translator pluginfrom the ACS policy file and pushed directly to Aviatrix from the CI/CD pipeline. CoPilot logs every allowed and denied connection with source identity, destination fully qualified domain name, and timestamp: continuous audit evidence, not snapshots. The plugin coverage report documents exactly which rules are enforced at the network today and which require the upcoming Guardrail Profile feature, so the control's scope is recorded, not assumed.

Questions Worth Asking

- How many distinct AI agent frameworks are running across your business units today, and do they share one security policy or several partial ones?
- If a vendor-shipped or compromised agent ignored its governance policy right now, what in your current stack would stop it from reaching an unauthorized MCP server or LLM API?
- Can you produce one audit trail of every tool endpoint your agents reach – across every framework, cloud, and runtime – or does each platform log separately?
- How are you currently proving to auditors that AI agent egress is controlled, not just described in policy?

What's Included

Component	Description
ACS-to-DCF Plugin	Open-source command-line tool that compiles a Microsoft Agent Control Specification policy file into Aviatrix DCF SmartGroups, WebGroups, allow/deny rules, and inline pattern guards. Emits a coverage report and applies policy via the REST API or Terraform provider.
DCF Egress Policy Templates	Pre-built WebGroup profiles for common MCP server and LLM endpoint sets, with a default-deny base policy included.
Inline Pattern Guards	ACS input and output validation patterns with semantic detection and redaction arrive with the Guardrail Profile feature, targeted for late summer 2026.
CoPilot Dashboard	Agent traffic visibility with every allowed and denied connection logged by source identity and destination fully qualified domain name, correlated with ACS policy events by session identifier.

Get Started

Request a 30-minute architecture review

We walk through the agent egress control model in your environment, inventory the agent frameworks and runtimes currently deployed, and surface the tool-call paths worth bringing under explicit policy.

About Aviatrix

Aviatrix® is pioneering the Cloud Native Security Fabric – the architecture the Containment Era requires. The Cloud Native Security Fabric governs every workload communication path across every cloud, every VPC, every Kubernetes cluster, and every serverless function, from a single policy plane. One rule. Universal propagation. Enforced at the workload, not at a chokepoint. Trusted by more than 500 of the world's leading enterprises. For more information, visit aviatrix.ai.