AVIATRIX®

# Contain the Blast Radius of Your AI Agents

## Why the age of autonomous AI needs containment, not faster patching.

Every company is racing to ship AI agents: software that acts on its own, reaches out to tools, models, and data, and makes its own connections without a person in the loop. Each agent is a new and highly privileged identity on your network, and they are being built and deployed faster than any security team can review them. That is the opportunity. It is also the exposure.

## AI Workloads are Where the Risk Concentrates

Three things make AI workloads the target:

**01**

They are **short-lived**, so defenses that depend on installing software on each machine cannot keep up.

**02**

They are **highly privileged**, holding the credentials and the reach to touch sensitive systems.

**03**

They are **shipped fast**, outrunning security review.

Put those together and risk concentrates on exactly the workloads moving quickest. This is what we call the AI Workload Bullseye, and attackers understand it better than most defenders do.

## The Math has Already Decided

We have been told the answer is to patch faster and detect sooner. The numbers say that race is already lost. Across thousands of organizations, remediation effort has risen **6.5 times** while outcomes have gotten worse. Exploits now arrive **about 7 days** before a fix is even available. And **82% of intrusions** ride valid credentials through legitimate channels. When the attacker signs in with real credentials through an approved path, there is no anomaly to catch. Detection and patching, on their own, cannot save you from that.

## Ask a Better Question

When something is compromised, how far can it reach? That distance is the Blast Radius. Bound it, and a compromise stays a contained incident instead of becoming a catastrophic breach.

> *When prevention fails and detection is too slow, containment decides whether the incident becomes a catastrophic breach.*

## What is Containment?

Containment is the architectural enforcement of one simple rule at every workload: it may talk only to what it is explicitly allowed to talk to, on every path available to it, whether or not anything has been detected. It is not a dashboard that tells you, after the fact, how exposed you were. It is enforcement, in the network path, at the identity of the workload itself.

## Why Does Containment Fit Agents in Particular?

An AI agent is powerful because of its reach: the tools, models, and data it is allowed to call. Containment governs exactly that reach, by the agent's own identity, on every path it has, including the paths that never pass through a central inspection point. It holds whether the agent runs as a managed cloud service, as a container on Kubernetes, or as a serverless function, and it does so without installing anything on the workload itself. The agent does not even have to know it is there.

## You Do Not Have to Choose Between Speed and Safety

This is the part that matters for the teams building AI and the teams securing it. You do not have to slow your AI teams down to contain them, and you do not have to rebuild your network to get there. The enforcement runs on infrastructure you already operate, it starts in a watch-only mode and switches on with a single change, and it reverses just as fast. Security gets control without becoming the bottleneck. The builders keep their speed.

**Contain the Blast Radius of your AI agents, on the infrastructure you already run. That is the work of this era.**

**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.