

The Security and Authorization Layer for Agentic AI

Extending Zero Trust to the Agent Action Layer

The Problem We Solve

Agentic AI has moved from experimentation into enterprise workflows. What began as tools that generate reasoning and information is evolving into autonomous agents that can access systems, use credentials, interact with data and execute tasks that affect the world. As organizations integrate agents into business processes, they face a new challenge: controlling what those agents are allowed to do, now they can act.

This shift fundamentally changes the nature of enterprise risk. An agent can be manipulated by an attacker, influenced by the content it consumes, or simply reason its way towards a goal, taking actions nobody anticipated.

Existing approaches to identity and access management, runtime enforcement and AI

“

“Harnessing the power of agentic AI requires a fundamentally new security architecture - one that binds identity, authorization, and action into a single operation at the moment of execution.”

Dr Peter Vincent
Co-founder and CTO, Outerlimit

governance cannot provide the provable observability and effective control required to scale agentic deployments safely.

As a result, enterprise adoption remains constrained by caution and a lack of trust rather than capability.

OUTERLIMIT - A NEW PARADIGM TO SECURE AGENTIC AI

Recognizing the need for a fundamentally new security architecture, Outerlimit provides the world's first decentralized security and authorization layer designed specifically for securing agentic AI.

Extending Zero Trust to the agent action layer for the first time, Outerlimit uses a distributed architecture and cryptographic enforcement at the moment of tool execution. Rather than storing credentials, keys or secrets in a single location, they are fragmented across the agent ecosystem and can only be reconstructed at the moment a tool is invoked. Decryption occurs only when the required identity, policy and execution context are present and verified, ensuring every action is authorized before it executes. This moves Zero Trust beyond verifying who is acting and what it can access to deterministically control the action an agent is permitted to take.

BACKED BY


evo/ution
EQUITY PARTNERS

AlbionVC

 **Crane**

The Pathway to Deterministic Enforcement

Outerlimit provides a structured pathway to enforcement, meeting customers where they are today. The platform enables enterprises to continuously discover and observe agentic tool use, quantify risk, and then enforce deterministic controls at the point of tool execution:



What makes us different

01 Deterministic, not probabilistic

Every agent action is on-policy or off-policy, decided at the point of execution.

02 Provable observability

Tool-call-level granularity provides visibility across all agent frameworks and all action surfaces, creating one common language.

03 Zero credential exposure

A credential is reconstructed only at the enforcement point; nothing standing for an agent or intermediary to hold or replay.

04 Provable multi-hop chain integrity

Authorization is bound to the composite identity of the full agent chain, to which every agent cryptographically contributes, so a compromised or influenced intermediary cannot extend authority down the chain.

05 In-tenant deployment

Deployed as a virtual appliance inside a customer's environment.

06 No central dependency

Enforcement is local to each point of action, with no latency round-trip and no single point of failure.

The Founding Team

Outerlimit was founded by cybersecurity entrepreneurs Tony Pepper and Neil Larkins, who built Egress Software, the category leader in cloud email security which sold to Vista Equity Partners backed KnowBe4 for \$400m in 2024.

Tony and Neil have partnered with Dr Peter Vincent, a theoretical neuroscientist with a PhD from the Sainsbury Wellcome Centre and the Gatsby Computational Neuroscience Unit at UCL.