PARLON MCP & AGENT ACTIVITY SOLUTION

AI agents are the newest thing touching your infrastructure, and the least accountable.

The Parlon MCP & Agent Activity Solution gives security and platform teams a live, audited record of what every AI agent reads, tied to identity and scoped by permission, inside the boundary you already control.

THE PROBLEM

AI agents and MCP servers are showing up across the infrastructure stack: querying telemetry, correlating incidents, drafting recommendations. Most of them run on borrowed credentials, with no record of what they saw, no scope beyond "everything the human could see," and no audit trail if something goes wrong.

The question a security team asks isn't "is the model good?" It's "what did it touch, under whose authority, and can we prove it?" Most teams can't answer that today, because agent activity isn't logged as agent activity at all — it's indistinguishable from a human logged in with the same account.



Why this matters for the AI buildout: this is the second wedge of the same thesis as GPU, PDU, and Network Path. AI now operates inside infrastructure that, in regulated or air-gapped environments, can't let telemetry, credentials, or write access leave the boundary. Governance has to be built into the substrate, not bolted on as a dashboard feature.

WHAT IT DOES

What the Parlon MCP & Agent Activity Solution does

ID

Every read, attributed

Every MCP query is tied to a known agent identity, not a shared service account, so “who looked at what” is never a reconstruction project.

SCOPE

Scoped by permission

Agents get read access to exactly what they're granted, nothing more. Scope is enforced at the platform, not assumed from the model's good behavior.

AUDIT

Full audit trail

Every read is logged: what was queried, when, by which agent, under which policy. Exportable for security review.

LIVE

Live supervision

A real-time view of what every agent is doing right now, not a report generated after the fact.

47.7%

WHY IT MATTERS

of enterprises already run AI training or inference workloads today. Most don't have a way to prove what those workloads, or the agents helping run them, actually touched.

Source: EMA, 2026

DEPLOYMENT & COVERAGE

One boundary, no exceptions.

Agent activity flows through the same Normalization Engine as every other infrastructure source, and deploys wherever Parlon does: SaaS, on-premises, hybrid, or fully air-gapped, so agent oversight never has to leave the boundary it's watching.

MCP (Model Context Protocol) **Read-only agent interface**

LEARN MORE