RUNTIME AUTHORIZATION FOR AI AGENTS

Lumen A2A Gateway

The policy decision point and policy enforcement point for agent intent

Overview

AI agents now initiate transactions against enterprise applications over MCP and A2A. Perimeter controls see the traffic; they do not understand the intent. Lumen A2A Gateway closes that gap as the runtime policy decision point and policy enforcement point for agent intent, evaluating every agent interaction against Smart Policies defined per PeerGroup before it reaches an MCP server or enterprise application.

The gateway discovers agents and resources across the environment, assigns each a risk score, and places them in PeerGroups for microsegmentation. Content flows through inline inspection over ICAP (RFC 3507), so existing security infrastructure participates in every decision. Agent identity is established with SPIFFE SVIDs or built-in cryptographic attestation with mTLS, and access federates into the identity stack enterprises already run: bearer tokens, Microsoft Entra agent blueprints, and OAuth 2.0 token exchange (RFC 8693).

Why it matters

Authorize intent, not just traffic

Every agent interaction is evaluated against Smart Policies per PeerGroup and allowed, denied, or constrained at runtime.

See every agent and resource

Continuous discovery with risk scores turns shadow agents into governed, microsegmented PeerGroups.

Keep the security stack in the loop

Inline inspection over ICAP (RFC 3507) lets the existing inspection infrastructure rule on agent content.

Prove it to auditors

First-class interaction and PeerGroup access logs record every decision for SIEM, compliance, and forensics.

Capabilities

Agent and resource discovery

Continuous discovery of AI agents, MCP servers, and resources across the environment, each carrying a risk score that informs policy.

PeerGroup microsegmentation

Agents and resources are grouped into PeerGroups, the unit of segmentation and policy. Interactions outside a PeerGroup's Smart Policies never reach the application.

Smart Policies (PDP and PEP)

Intent-level authorization defined per PeerGroup and enforced inline, grounded in OWASP agentic threats: intent breaking, tool misuse, and goal drift.

Inline inspection over ICAP

Standards-based ICAP (RFC 3507) steers agent payloads through existing inspection services, so content rulings happen inside the flow, not after it.

Agent identity and mTLS

SPIFFE SVIDs or built-in cryptographic attestation give every agent a verifiable identity; sessions are mutually authenticated with mTLS.

Federated IAM

Bearer tokens, Microsoft Entra agent blueprints, and OAuth 2.0 token exchange (RFC 8693) connect agent access to the enterprise identity stack, unchanged.

Architecture

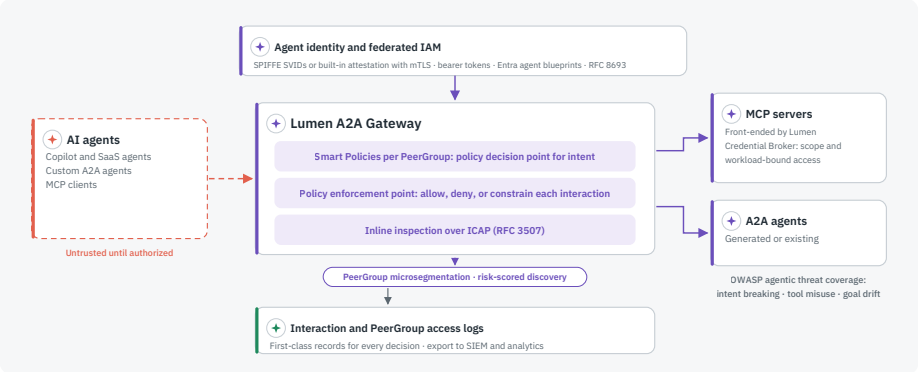

Lumen A2A Gateway authorizes agent intent per PeerGroup before any call reaches an MCP server or A2A agent.

How it works

1. An agent's request enters the gateway through the steering method that fits the topology: ICAP from an existing proxy, policy-based routing for east-west traffic, or a VIP with destination NAT.
2. The gateway establishes the agent's identity, resolves its PeerGroup, and evaluates the requested interaction against Smart Policies: the policy decision.
3. Approved content passes through inline ICAP inspection, then the gateway enforces the verdict, allowing, denying, or constraining the call to the MCP server or enterprise application: the policy enforcement.
4. Every interaction and PeerGroup access is logged as a first-class record and exported to SIEM and analytics.

Deployment

Deploys inline in the enterprise environment with three steering options: ICAP from an existing proxy or firewall, policy-based routing for east-west agent traffic, and VIP with destination NAT for published endpoints.

Works alongside the security and identity infrastructure already in place. Pairs with Lumen POET, which generates precise MCP servers and A2A agents from the API specifications enterprises already have.

Technical specifications

Function	Policy decision point (PDP) and policy enforcement point (PEP) for agent intent
Policy model	Smart Policies defined per PeerGroup
Discovery	Agents and resources, each with a risk score
Segmentation	PeerGroup microsegmentation
Inspection	Inline over ICAP (RFC 3507)
Agent identity	SPIFFE SVIDs or built-in cryptographic attestation; mTLS
IAM federation	Bearer tokens; Microsoft Entra agent blueprints; OAuth 2.0 token exchange (RFC 8693)
Traffic steering	ICAP; policy-based routing for east-west traffic; VIP with destination NAT
Protocols governed	MCP (Model Context Protocol), A2A (Agent2Agent)
Threat coverage	OWASP agentic threats: intent breaking, tool misuse, goal drift
Logging	First-class interaction and PeerGroup access logs; SIEM export

Standards and interoperability

- **MCP:** Model Context Protocol, modelcontextprotocol.io
- **A2A:** Agent2Agent protocol, a Linux Foundation project
- **ICAP:** RFC 3507, Internet Content Adaptation Protocol
- **OAuth 2.0 token exchange:** RFC 8693
- **SPIFFE:** CNCF graduated workload identity framework
- **Zero trust:** Aligned with NIST SP 800-207 and SP 800-207A
- **Threat model:** OWASP GenAI agentic threats

THE BACKFLIPT DIFFERENCE

No agent is re-coded. No application authorization is reworked. Enterprises scale AI on the identity and security investments they already have.

About Backflipt

Backflipt (Xenovus, Inc.) builds the Lumen product line for runtime authorization of AI agents. Lumen A2A Gateway authorizes agent intent, Lumen POET generates agent-ready MCP servers and A2A agents from the API specifications enterprises already have, and the Lumen POET Credential Broker keeps standing secrets in the vault and out of agents. www.backflipt.com · reachus@backflipt.com