

Lumen POET

Generate precise MCP servers and A2A agents from the APIs enterprises already have, no code

— Overview

Enterprise APIs were never designed for AI. Operation names are cryptic, descriptions are thin or missing, and a single specification can expose hundreds of operations an agent should never see. Point an LLM at a raw specification and the result is noisy tools, wrong calls, and unpredictable behavior.

Lumen POET turns those same specifications into agent-ready building blocks. It imports OpenAPI, Swagger, or RAML, gives administrators full flexibility to edit and curate the schema for AI, and then generates precise MCP servers and complete A2A agents with out-of-box prompts, all without writing code and without changing the underlying API. Generated services deploy with Helm and Terraform inside the customer VPC and coexist with the API gateways already in production.

— Why it matters

Curate the schema for AI

Edit operations, parameters, and descriptions before generation. The API stays untouched; the agent-facing schema becomes precise and intelligible.

MCP servers that are precise

Generated MCP servers expose only the curated operations, so agents see exactly the tools they should use and nothing else.

A2A agents, ready on day one

Auto-generated A2A agents ship with out-of-box prompts and agent cards. Run them as generated or edit every prompt.

No code, no rework

Generation is spec-driven end to end. No agent is re-coded and no application is modified.

— Capabilities

Specification import

Parses OpenAPI, Swagger, and RAML, including operations, parameters, schemas, and authentication definitions.

Schema curation for AI

APIs were never designed for AI, so POET makes them agent-ready: select and rename operations, prune parameters, rewrite descriptions so an agent understands what each tool does, and scope what generation exposes. The underlying API is never modified.

MCP server generation

Generates precise MCP servers scoped to the curated schema, ready for any MCP client. No code is written.

A2A agent generation

Generates complete A2A agents with out-of-box prompts and agent cards. Agents run immediately and every prompt remains editable as behavior is tuned.


VPC deployment

Helm charts and Terraform modules deploy generated services inside the customer VPC. Nothing leaves the enterprise boundary.

API gateway coexistence

Generated services sit alongside Kong, Apigee, AWS API Gateway, and MuleSoft. Existing API management investments stay in place.

Architecture



Lumen POET imports the specification, lets administrators curate the schema for AI, and generates precise MCP servers and A2A agents with out-of-box prompts.

How it works

- 1 Import the specification. POET parses OpenAPI, Swagger, or RAML and inventories every operation, parameter, and auth definition.
- 2 Curate the schema for AI. Administrators edit the agent-facing schema: keep the operations agents need, drop the rest, and rewrite names and descriptions until they read the way an agent reasons. The API itself is never changed.
- 3 Generate. POET produces a precise MCP server, a complete A2A agent with out-of-box prompts and an agent card, or both, from the curated schema. No code.
- 4 Deploy and govern. Helm and Terraform place the services in the customer VPC. Add the Lumen POET Credential Broker for credential-level PDP and PEP, and front the deployment with Lumen A2A Gateway for intent-level authorization.

Deployment

Generated MCP servers and A2A agents deploy with Helm and Terraform inside the customer VPC, alongside Kong, Apigee, AWS API Gateway, and MuleSoft.

The Lumen POET Credential Broker add-on front-ends generated services with credential-level policy, and Lumen A2A Gateway adds intent-level runtime authorization in front of the whole estate.

Standards and interoperability

MCP: Model Context Protocol, modelcontextprotocol.io

A2A: Agent2Agent protocol, a Linux Foundation project

API specifications: OpenAPI, Swagger, RAML

Packaging: Helm, Terraform

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

Backflipt, Lumen, Lumen A2A Gateway, and Lumen POET are trademarks of Xenovus, Inc. All other trademarks are the property of their respective owners. © 2026 Xenovus, Inc. All rights reserved.