

Standardizing Agentic Workflows with a Zero-Trust Platform

AT A GLANCE

Challenges

- Fragmented AI agent solutions
- No company-wide standards
- Zero-trust authentication requirement
- Platform-agnostic design needed
- Shared tooling across teams

Methods

- Collaborative architecture design
- Zero-trust pass-through authentication
- YAML-based agent definitions
- Kubernetes prototype with PostgreSQL
- Platform-owned lifecycle APIs

OBJECTIVES

The client needed to move from fragmented, team-by-team AI agent development to a unified platform that standardizes how agentic workflows are built and run. The platform had to provide an easy onramp for teams, enforce company-wide standards, and offer shared tooling, all without locking the organization into a single model provider or compromising individual team security boundaries.

IntelliTect brought embedded expertise to accelerate the design and validation of this platform. By collaborating directly with the client's technical leads over a focused 6-week engagement, IntelliTect architected a zero-trust, model-agnostic solution and delivered a working prototype that confirmed the viability of the approach, giving the client a clear foundation for productization.

SOLUTIONS

IntelliTect designed and prototyped a shared platform where agent definitions are authored in YAML and managed through platform-owned API endpoints. The platform uses a zero-trust pass-through authentication model, holding no security keys for individual teams. Built on Kubernetes with PostgreSQL for durability, the architecture supports MCP tools (team-owned and shared) and is model-agnostic by design.

OUTCOMES



Architecture Validated

Working prototype validated the zero-trust, platform-agnostic architectural approach.



Ready for Production

Clear path to productization established.



Realistic Timeline

Delivered in 6 weeks by 2 IntelliTect engineers.

