



CASE STUDY

How AI Turned Operational Data into Self-Service Intelligence for GasOps



Project Overview

Client GasOps

Solution

- AI-powered natural language reporting and operational intelligence layer

Industry

- Gas Utility Operations Software

Timeline

- 10–14 weeks
(4–6 weeks data integration + 6–8 weeks AI development; ongoing post-launch iteration)

Tech Stack

- Microsoft Fabric · AI/LLM-based NLP · REST API integrations (.NET) · Data pipelines · Frontend chat interface

Team Composition

- Backend Developers
- Data Engineers
- AI/ML Engineers
- Frontend Developers
- QA



Client Background

GasOps supports utility operations through a platform that spans multiple workflows and data domains, including Weld, Transmission Work Orders, OQ, Routesheets, and Inspections. These modules generate large volumes of operational and compliance data that are essential for tracking field activity, monitoring performance, and maintaining audit readiness.

As the platform expanded, so did the challenge of turning that data into timely, usable insight. Teams had the data, but not an efficient way to access it dynamically.

The Situation

GasOps builds enterprise software for utility operations; a platform that spans weld management, transmission work orders, operator qualifications, routesheets, inspections, and compliance workflows. It's mission-critical infrastructure for the people who keep gas networks running.

As the platform grew, so did the volume of operational and compliance data flowing through it. That data was there. The problem was what it took to access it.

The Problem

Getting useful information out of GasOps required manual effort and technical involvement. Users—operations managers, compliance officers, inspectors, field supervisors, contractors—submitted report requests and waited.

Static dashboards existed, but they were built around predefined queries that couldn't adapt to the dynamic, unpredictable nature of real operational questions.

Cross-module analysis was particularly difficult. Insights that spanned weld activity, OQ compliance, and transmission work orders required someone technical to pull and stitch together data from different parts of the platform.

The engineering team fielded these requests repeatedly, consuming capacity that should have been spent building.



Why it mattered

The cost of doing nothing wasn't an administrative inconvenience. It was operational friction with direct consequences.

Delayed reporting slowed decisions. Manual workflows consumed engineering resources that belonged elsewhere. Compliance insights surfaced too slowly for teams that needed them in the moment. And a large volume of available data remained effectively locked; present in the system, but inaccessible to the people who needed it most.

GasOps had the data. What it lacked was a scalable way to turn that data into day-to-day decision support.





The Solution

AthenaWorks built an AI-powered operational intelligence layer that puts data access directly in the hands of users; no technical intermediaries required.

The solution enables natural language querying across GasOps modules.

Users ask questions in plain language:

- "Show pending weld inspections."
- "Which contractors have expired OQs?"

The system interprets the question, retrieves the relevant data, and returns an accurate, context-aware answer; instantly.

The architecture is built on Microsoft Fabric, which serves as the centralized operational data platform. Fabric consolidates and structures data from across GasOps modules — weld operations, OQ compliance, transmission work orders, inspections — into unified, AI-ready datasets. REST API integrations (.NET) connect live operational data to the AI layer, ensuring responses draw from current system data rather than static snapshots.

The AI layer handles natural language interpretation, structured query orchestration, and intelligent summarization. A conversational frontend interface makes the experience approachable for both technical and non-technical users.

AI was the right choice because operational users don't ask the same limited set of questions repeatedly. They need to explore dynamic, context-dependent questions across modules and workflows.

The system needed to interpret intent, understand domain-specific terminology—and return reliable answers at speed.

The rollout was phased:

- Data integration and pipeline setup in approximately 4–6 weeks
- Followed by AI agent development over 6–8 weeks

With ongoing iteration post-launch. The team deliberately scoped early phases around core reporting use cases before expanding capabilities based on real user behavior; the right tradeoff between speed to value and long-term quality.



The Results

After launch, GasOps users could interact with operational data directly; no report requests, no waiting, no technical support required for day-to-day insight needs.

Outcomes:

- Report generation time reduced by approximately 70–80%; work that previously took 30–60 minutes of manual effort now completes in under 5 minutes
- Manual reporting and ad hoc support requests to technical teams reduced by approximately 50–60%
- Faster decision-making for operations and compliance teams across Weld, OQ, and Inspections
- Improved data utilization — information that was previously locked behind manual processes became accessible to a broader set of users
- Stronger compliance monitoring and improved cross-module visibility
- High adoption across both technical and non-technical audiences, driven by the low-barrier chat interface

Beyond the immediate gains, the solution positioned GasOps for the next phase of intelligence. The AI layer is now core infrastructure — a foundation for advanced analytics and future predictive capabilities that weren't possible with the previous reporting model.

The Takeaway

GasOps didn't need more data. It needed a better way to use it. Operational intelligence shouldn't require a support ticket.

Contact Us

athenaworks.com

