Ultimarii

CASE STUDY

Enabling AI Adoption in a Regulated Utility Without Compromising Trust



The Challenge

AI Interest Meets Non-Negotiable Security Standards

A Canadian electricity transmission and distribution utility operates under constant scrutiny from provincial and federal regulators. Its legal and regulatory teams manage highly sensitive materials every day—rate applications, environmental filings, capital project approvals, and commercial agreements where data governance is not a preference but a requirement.

The organization saw clear potential for AI to reduce research time, improve consistency, and support regulatory submissions. But adoption hinged on one condition: **AI had to fit inside their existing security and compliance framework.**

The utility's requirements were explicit and uncompromising:

- All proprietary documents hosted privately
- Data residency within their designated Azure region
- Strict separation between public regulatory materials and internal content
- No model training or reuse of private data
- Full alignment with existing audit, identity, and access controls

Most AI platforms failed this test. Off-the-shelf tools required data to leave the organization's environment, while custom deployments demanded months of engineering effort and long approval cycles. For a utility where security approval is a prerequisite—not a downstream task—both options stalled progress.

Ultimarii's Approach ↘

Designing AI Around the Trust Perimeter

Ultimarii did not ask the utility to compromise.

Instead, Ultimarii worked directly with the utility's IT and cybersecurity teams to design a deployment that aligned with how the organization already manages risk. The goal was simple: **enable AI capabilities without changing where data lives or how it is governed.**

Private libraries were provisioned directly within the utility's own Azure environment. Proprietary documents never left their infrastructure, while Ultimarii's platform—retrieval, citations, playbooks, and curated public regulatory libraries—remained fully functional on top of that foundation.

This was configuration, not customization. Ultimarii mapped its architecture to the utility's existing controls, ensuring adoption could proceed without lengthy redesigns or exceptions.

The deployment included:

- Full tenant isolation for private libraries
- Direct Azure file storage integration
- Zero replication or training on private data

Security validation and implementation were completed in days, not weeks or months due to the flexible nature of the solution.

The Solution

Making Governed AI Usable for Day-to-Day Work

Once the environment was approved, teams moved quickly from evaluation to use.

With trust established at the infrastructure level, legal and regulatory staff were able to apply AI to workflows that had previously been off-limits:

- Drafting and reviewing regulatory briefs using private precedents
- Comparing submissions across historical rate cases
- Analyzing information requests and preparing structured responses
- Reviewing contracts against internal standards and prior agreements

Because Ultimarii supported both **private internal libraries and curated public regulatory content**, teams could work across the full context of a proceeding—without manual document switching or governance risk.

The Impact

Adoption Enabled by Assurance

The most significant outcome was not speed alone, but certainty.

- **Security approval achieved without exception requests:** AI adoption moved forward inside existing compliance frameworks.
- **Time to first use measured in weeks:** The configurable architecture was set up entirely though remote coordination within a few days to enable smooth onboarding.
- **Expanded AI use cases:** Teams were able to apply AI to sensitive regulatory and legal work, not just public research.
- **Reduced internal friction:** Legal, regulatory, and IT teams aligned on a single, trusted deployment model.



The Bottom Line

Trust as the Enabler of Scale

For this utility, AI adoption succeeded because it respected existing risk management practices. Value followed assurance, not the other way around.

For Ultimarii, this engagement demonstrates a core principle: **AI scales in regulated industries only when it is designed around trust, not layered on top of it**. By aligning to customers' infrastructure, governance, and operating reality, Ultimarii enables organizations to move from cautious interest to real, defensible adoption.