

Agentic AI Implementation for Energy Project Risk Review

BUSINESS PROBLEM

NextEra Energy Resources' Discover platform identifies promising energy project sites but could benefit from more early-stage, high-quality risk insights, especially for community engagement. Early-stage project advancement without sufficient risk review can lead to later-stage challenges, including project attrition and increased development costs. Current SME-driven review processes require significant time investment each year and may introduce scheduling bottlenecks, as well as variability in evaluation outcomes. There is an opportunity to enhance early risk assessment by improving speed, scalability, and consistency in review processes.

DATA SOURCES


1. NextEra Project Database – structured attributes and outcomes for thousands of past and present projects
2. SME Risk Reviews – unstructured text comments and qualitative ratings
3. Web Data (Tavily) – public sources on ordinances, sentiment, and opposition signals
4. SME surveys – feedback surveys from SMEs

Data Types and Format

Boolean, Categorical (nominal), Numeric, Structured, Text, Unstructured

APPROACH

Built an agentic AI system integrated into Discover using Amazon Bedrock and Claude models, orchestrated via a modular agent framework with web search (Tavily) and retrieval capabilities. Combined ML analysis (logistic regression, XGBoost, TF-IDF on SME text) to identify key risk drivers and inform prompt design. The agent performs autonomous search, extracts relevant data, and generates structured risk summaries, with outputs precomputed and stored for low -latency retrieval in the UI.



IMPACT

The AI agent significantly improves speed, scalability, and consistency of early-stage risk assessment. By providing instantaneous, precomputed risk insights within Discover, it eliminates the lengthy SME reviews for initial screening, enabling near real-time decision-making.

The solution directly reduces SME workload by partially automating community engagement analysis, targeting a portion of the substantial SME workload associated with reviews.

From a cost perspective, earlier identification of fatal flaws can prevent projects from advancing to later stages where sunk costs are significant. Avoiding even a small number of failed projects can yield substantial savings and improve capital allocation efficiency.

Qualitatively, the system improves consistency of risk assessments by standardizing outputs and grounding them in structured prompts and external data, reducing subjectivity across SMEs. It also embeds actionable insights directly at the point of site selection.

Finally, the project establishes a scalable architecture and evaluation framework for enterprise AI deployment, positioning the organization to expand agentic AI across additional risk domains (e.g., environmental, permitting), compounding long-term operational and financial benefits.

 DRIVERS	Elevated project attrition and sunk costs highlighted opportunities to strengthen early risk identification. SME-driven reviews created bottlenecks. Leadership aligned on the need for faster, more scalable, and consistent decision support, while advancing enterprise AI adoption.
 BARRIERS	Limited labeled data prevented model fine-tuning; structured data had weak predictive power. Integration into AWS-based systems constrained tooling choices. Ensuring trust in AI outputs and aligning with SME workflows were critical adoption challenges.
 ACTIONS	Analyzed historical project data with ML and NLP to identify risk drivers. Selected community engagement as a high-impact use case. Co-designed prompts with SMEs, built an agentic AI system using Bedrock + web search tools, and integrated outputs into Discover. Developed a multi-metric evaluation framework and iterated based on SME feedback.
 INNOVATION	Combines agentic AI with retrieval, web search, and SME-informed prompt design to replicate qualitative risk assessment. Creates a dynamic, tool-using agent embedded in production workflows. Introduces a hybrid architecture with precomputed outputs for scalability and low-latency.
 IMPROVEMENT	Provides real-time, high-quality insights on community engagement risks enhancing visibility beyond prior capabilities. Reduces SME workload through partial automation. Strengthens consistency and accessibility of risk insights, enabling earlier identification of fatal flaws and more informed project advancement decisions.
 BEST PRACTICES	Start with a narrow, high-impact use case; co-design with end users to drive adoption. Use modular architectures and evaluate across multiple dimensions (accuracy, reliability, usability). Prioritize augmentation over replacement to build trust and enable gradual AI integration.
 OTHER APPLICATIONS	Extend to other SME domains (environmental, permitting, transmission) and infrastructure sectors (real estate, mining, transportation). Applicable to any workflow requiring early-stage risk screening using unstructured data and external signals.