

AI-Powered Knowledge Platform For a Wind Turbine Manufacturer

How developing a conversational AI knowledge platform with offline access and peer-to-peer expertise sharing enabled a European wind turbine manufacturer to turn post-sale field support into a **product differentiator** and reduce support queries by **37%**.



Business challenge

When a turbine stops or triggers an alarm miles from the nearest office, the repair depends on how fast the technician on site can find the right answer. For field workers servicing wind turbines offshore and over sprawling onshore locations, **that process is riddled with friction.**

- **No single source of truth.** Maintenance manuals, monitoring guides, technical bulletins, and troubleshooting references exist, but are scattered across multiple portals and PDF libraries with no unified entry point. A technician looking for a specific calibration procedure might find it in three clicks or thirty, depending on which system they open first and whether they guess the right keywords.
- **Documentation built for desks, used on turbines.** The materials are authored for screen-and-keyboard reading, but real-life wind turbine servicing conditions look nothing like that. Technicians access documents on phones or tablets, often one-handed, in rain or wind. Dense PDF layouts with no mobile optimization turn a two-minute lookup into a scrolling exercise.
- **Peer knowledge without a peer network.** When documentation falls short, experienced technicians fill the gap. But that expertise is shared through phone calls, ad hoc messages, or chance conversations rather than through a structured channel. Two technicians in the North Sea and Andalusia might independently solve the same fault, weeks apart, with no way to learn from each other.

Our client, a major European wind turbine manufacturer, saw these hurdles as a product opportunity. As their relationship with every buyer extends into years of customer support, **an application that would consolidate the company's product maintenance knowledge into a single, searchable hub could not only ease post-sale support overhead but become a selling point in its own right.**

They approached Instinctools with a **vision for a companion app built around conversational AI: a tool, offered alongside the turbines themselves, giving field technicians a way to describe a problem in plain language and get an actionable, source-verified answer on the spot, with clear traceability back to the original material.**

Solution

Some opportunities can't wait for paperwork. While the project was still in the proposal stage, Instinctools built a **working prototype** that gave the client an early view of the product vision in action. It allowed their team to test assumptions, gather internal feedback, and validate the concept before committing to a vendor, budget, or full-scale delivery plan.

01 Crafting a clickable prototype at the pre-sales stage

The client's original brief centered on a searchable knowledge base with a conversational AI interface. Field technicians would type a question in their native language, and the system would return a plain-language answer drawn from the manufacturer's documentation, with a citation trail pointing back to the source. Multilingualism was essential as the client's install base spans dozens of countries.

Our **AI developers** and **LUX/UI designers** translated the brief into a **clickable prototype covering the core user journey:**

- 📁 Browsing the document library
- 🗉 Querying the AI assistant
- 📄 66 99 Reviewing cited sources

What started as Instinctools' suggestion, by the end of the prototype review landed in the MVP scope, transforming a knowledge base app into a broader **AI-assisted field support platform grounded in source documentation and strengthened by a technician network.**

02 Laying the technical foundation for full-scale development

With the prototype approved and the scope expanded, Instinctools moved into architectural planning. Several technical decisions shaped the project from the start:

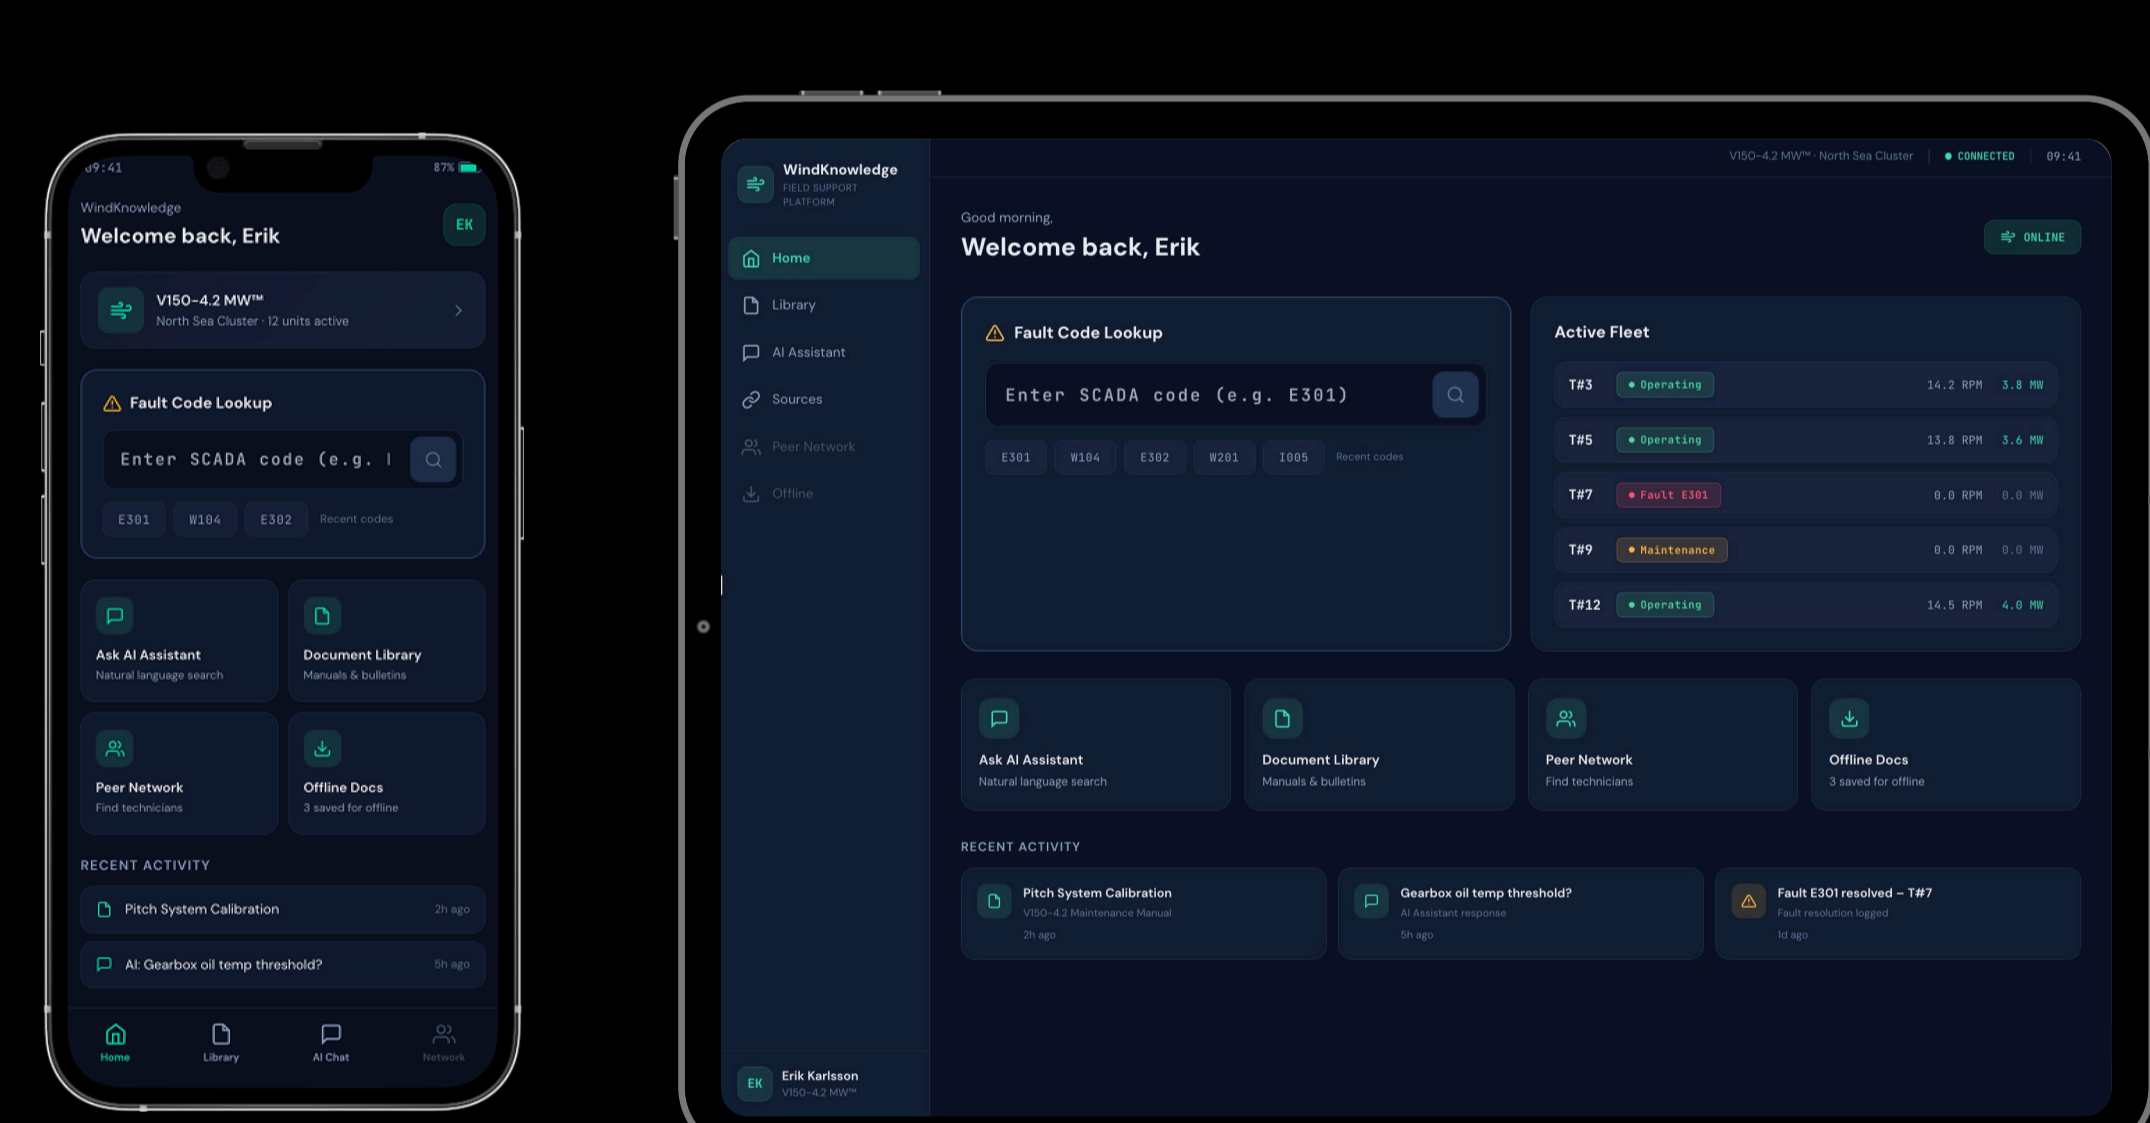
- **Single-vendor cloud and AI setup.** We hosted the platform on AWS, picking the eu-central-1 region in Frankfurt to keep data physically inside the EU. **AWS Bedrock** handles the AI layer, providing access to Anthropic's **Claude** within the same compliance boundary as the rest of the infrastructure.
- **Search architecture.** The accuracy of a conversational AI assistant is rooted in the quality of the retrieval layer feeding it. Instinctools chose **pgvector** to combine semantic search with automatic profile-based hard filters in a single query, so when a technician asks about pitch system calibration, the results are already scoped to a particular turbine model. For a knowledge base sized at around 10,000 content chunks, this vector similarity search keeps the infrastructure lean without cutting corners on retrieval accuracy.
- **Zero-signal mode.** Because technicians may work offshore or in low-connectivity environments, the app needed to stay useful even without a stable connection. We paired **Workbox** and **IndexedDB** to cache the app interface and store bookmarked documents locally within the app. Features that rely on live services, like AI search and community messaging, stay online-only.
- **Access control.** For authentication, Instinctools went with **Auth0**, starting with invite-only registration during the MVP, so the client could control early access before opening to self-service registration. Password policies, account lockout, email verification, and recovery flows were handled through Auth0's built-in identity-management capabilities.
- **Autotranslation layer.** The client's turbines operate across a multilingual EU market and beyond. The team chose **DeepL's API** for on-demand translation, favoring it over alternatives for its stronger handling of technical content in European languages.

03 Refining the prototype with field feedback

The prototype looked right on screen. To see if it'll be of any help for end users, we shared it with a group of roughly **20 early testers** the client gathered from its European customer base.

Several assumptions didn't survive first contact with the field, but caught early, they became opportunities for improvement before full-scale development:

- **Fault-code workflow.** The AI search handled natural language and fault codes equally well, but testers revealed that the vast majority of their queries start with a SCADA error code. This prompted us to update the home screen design by adding a dedicated fault-code entry field with auto-recognition, so the most common interaction pattern gets the shortest path.
- **Expert knowledge sharing priorities.** The prototype included a location-based map showing where other technicians were working, but users made it clear that proximity was not their top priority. They wanted to filter peers by turbine model and specialization first, and by location second. For a technician troubleshooting a specific gearbox issue, the right match was not the nearest colleague, but the one with relevant hands-on experience.
- **Minimum taps to target.** In the initial prototype, users had to open a document before saving it for offline access. Testers pointed out that before going offshore, they often need to save several guides at once. We redesigned the flow so documents could be saved directly from search results, without opening each one individually.



This feedback reshuffled the development backlog. We moved features like fault-code-driven search and single-tap document access from "could have" to "should have" status, ensuring the MVP reflected how technicians actually work, not just how the workflow looked in a prototype.

04 Developing and rolling out an MVP

The MVP was deliberately focused: one turbine product line, one content library, and an upscale from 20 to 100 technicians. The goal was to prove the full end-to-end experience with a larger, **entire user group** before replicating it across the client's entire equipment catalog.

Here's how the flow goes:

- 01 During onboarding, the early adopters select the turbine models they service
- 02 From that point on, the app recommends search results, document personalizations, and community connections, so that a technician working with a specific model sees materials and peers' contacts relevant to that model, noise-free
- 03 Every answer the AI assistant returns includes citations to the source document and section, so a technician can verify the guidance before acting on it.

This personalization engine sits on a knowledge base of hundreds of maintenance manuals, monitoring guides, technical bulletins, etc. Since it's updated once a month at most, a real-time content synchronization would have added unnecessary complexity. A monthly content ingestion pipeline was enough to keep the knowledge base current without overengineering the backend.

Once rolled out to the full 100-technician group, the app saw steady organic adoption.

Before

- Technical documentation scattered across multiple portals and PDF libraries
- Only basic, document-level search capabilities
- A multi-tap user journey to reach a single document, with no shortcut for the most common queries
- Peer expertise shared informally through phone calls and personal contacts, if at all

After

- A single AI-powered conversational app consolidating the wind turbines' maintenance library
- Advanced semantic search with optional hard filters across the whole knowledge base
- One tap to target and fault-code auto-recognition field on the home screen
- A structured peer network connecting technicians by turbine model, specialization, and relevant field experience

Business value

- Post-sale support transformed into a **product feature**
- **-37%** requests escalated to the client's customer support team
- **<30 seconds** to query to the client's source-verified answer
- **x2.5 faster** troubleshooting of field incidents

Multiplier effect

The pattern behind this platform – consolidating scattered documentation behind a conversational AI interface and pairing it with a peer network – is not limited to the energy sector. **Any industry where field professionals work with dense technical libraries under time pressure faces similar friction:** heavy machinery, medical devices, industrial automation, telecom infrastructure, and more. The infrastructure we built is modular to be useful for a different knowledge base and equipment catalog without being redesigned from scratch.



Do you have a similar project idea?

Contact us — and we'll estimate your project for free!

[CONTACT US](#)