

# Pre-Visit Intelligence — One-Page Case Study

---

AI-powered technician briefing that improves first-time fix rate.

---

<b>Role</b>	Solo PM + builder (0→1 MVP)
<b>Domain</b>	Appliance field service
<b>Stack</b>	FastAPI · SQLite · OpenAI · Next.js
<b>Live demo</b>	<a href="https://pre-visit-intelligence.vercel.app">pre-visit-intelligence.vercel.app</a>

---

## Problem

Appliance technicians often dispatch with only a model number and symptom. They arrive without the right parts, recall context, or a diagnostic plan — then schedule a return visit. Repeat truck rolls are a top cost driver in field service.

**North-star metric:** first-time fix rate (FTFR)

---

## Solution

Pre-Visit Intelligence turns sparse dispatch input into a scannable pre-visit briefing in one click: recalls, likely causes, recommended checks, tools, parts, duration, and an AI summary.

**Core bet:** local-first retrieval for trust, AI second for coverage — not "chat with a manual."

**Primary user:** Field technician — *"Before I drive, tell me what's wrong, what to bring, and if there's a recall."*

---

## What I Shipped

- Single-page intake + results dashboard with five sample scenarios
- One API endpoint (`POST /v1/cases/analyze`) returning structured briefing JSON
- SQLite knowledge base (recalls, manuals, failure modes, tools, parts, durations)
- Evidence-grounded AI summary + gap-fill for unknown appliance:symptom combos
- Graceful degradation without OpenAI (rules-based fallback)
- Deployed on Render + Vercel with automated tests

**Try it:** Whirlpool WRX735SDHZ · "ice maker not producing ice" → causes, parts (water inlet valve), checks, duration, and AI summary.

---

## Key Product Decisions

1. **Safety first** — Recalls surface above summary tiles; liability beats convenience.
2. **AI second, not AI only** — Prompts use evidence only; gap-fill carries explicit lower confidence (50–70%).

3. **Normalization before AI** — Symptom variants and model-family matching handle messy real-world intake.
  4. **One screen, one API** — Deferred auth, case history, and admin to ship a demoable MVP fast.
- 

## Outcomes

**Shipped:** Live demo, end-to-end pipeline, test coverage on representative cases.

**Next (pilot metrics):** Repeat visit rate delta, KB vs. AI fallback ratio, briefing helpfulness, generation time under 5s.

---

## What's Next

1. Pilot with one service org → measure FTR impact
  2. Persist cases + wire feedback loop
  3. Dispatch integration (ServiceTitan-class tools)
  4. Live recall feed + KB expansion
- 

## Lessons

- Field service products fail on **trust**, not model size.
  - **Sparse input** matching matters as much as summary UX.
  - A **live demo** beats a slide deck — employers can verify in two minutes.
- 

Built by Brian · [API docs](#)