

# AI Receptionist Pilot Spec

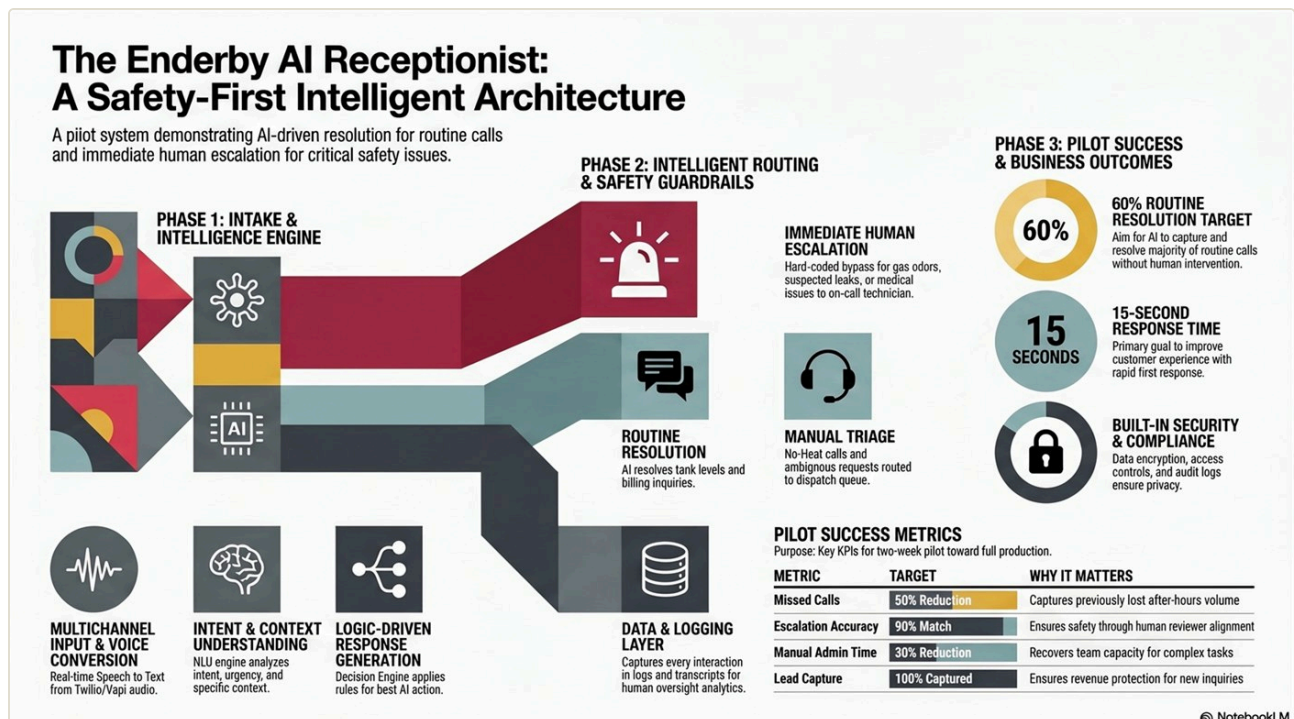
Enderby Gas - AI Operations Portfolio - Brianna Bates - briannabates1221@gmail.com

*A two-week pilot designed to validate - not deploy - a 24/7 voice agent at a propane distributor, where gas safety is the central concern.*

## Executive summary

Enderby handles a steady volume of inbound calls - delivery scheduling, tank monitoring, billing, emergency triage - all routed through a human dispatcher. That model works for complex situations and breaks down for everything else: after-hours calls go to voicemail, routine billing eats dispatcher time, and no-heat calls during cold snaps create handle-time spikes the team can't absorb.

The AI Receptionist is a voice-based frontline agent designed to answer calls 24/7, capture intent, resolve routine questions, and escalate urgent issues to the right human channel - without changing how Enderby fulfills service.



*The pilot architecture in three phases - multichannel intake and intent understanding, intelligent routing with safety guardrails, and the measurable outcomes that gate a production decision.*

## The business problem

---

Enderby's current call-handling model is fully human-driven. It creates five operational gaps:

- **Missed calls.** After-hours calls roll to voicemail; urgent issues wait until morning, creating next-day spikes.
- **Lost leads.** New-customer inquiries after hours or during peak hold times go cold before a CSR can call back.
- **Slow response time.** Peak-season hold times and repeated transfers create a perception of poor service.
- **Manual scheduling burden.** Dispatchers spend hours per week on calls that require no judgment.
- **No consistent documentation.** Call notes live in heads and notepads, so the next CSR starts over.

## Pilot objectives

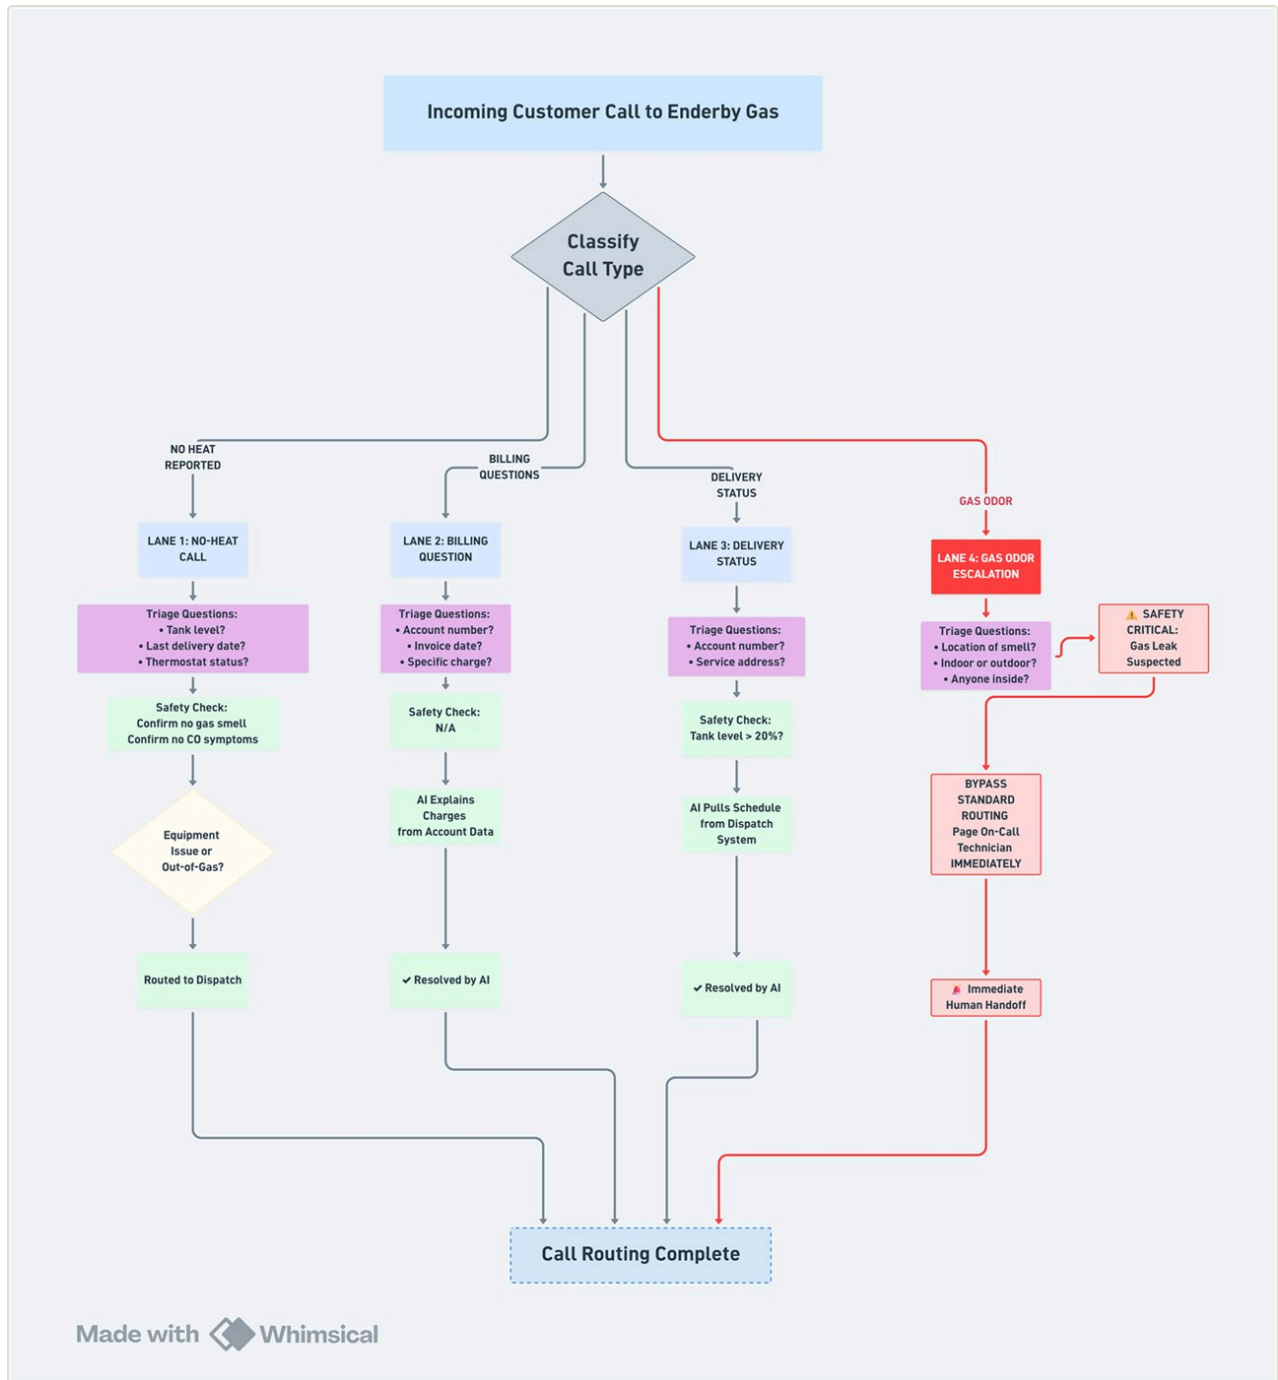
---

Designed to **validate, not deploy**. It must prove an AI agent can resolve at least **60% of routine calls** without a human, escalate emergencies with at least **90% accuracy** against human-reviewer judgment, and cut after-hours voicemail volume by at least **half**.

## Workflow - six steps on every call

---

1. **Customer calls.** A standard greeting identifies the agent as AI and offers immediate human handoff.
2. **AI answers.** Twilio/Vapi streams audio to the model, which asks an open-ended intent question.
3. **AI qualifies intent** - billing, delivery, tank level, new customer, or emergency. One clarifying question, then escalate if unclear.
4. **AI books, routes, or summarizes** using the NotebookLM knowledge base (Project 2 SOPs).
5. **Business gets a clean follow-up** - a structured summary to the dispatcher queue; emergencies trigger an immediate notification.
6. **Caller gets confirmation** of the next step.



The routing logic behind the guardrails: every call is classified into a lane with its own triage questions and safety check - and a gas-odor signal (red) bypasses standard routing entirely for an immediate human handoff.

## System architecture

LAYER	TOOL
Voice intake	Twilio Voice / Vapi
Reasoning	Claude / GPT-4o
Knowledge base	NotebookLM (Project 2 SOPs)
Scheduling	Google Calendar (API)
Lead tracker / CRM handoff	Google Sheet (Cargas API in production)
Follow-up	Twilio SMS + Gmail
Transcript storage	Google Drive
Reporting	Looker Studio dashboard

## Risks & guardrails

- **Emergency calls route to a human.** Any gas odor, suspected leak, CO alarm, no-heat with a vulnerable resident, or critically low tank triggers immediate handoff - no diagnostics, no holding.
- **Ambiguity escalates after one clarification** rather than letting the agent keep probing.
- **Consent & recording disclosure** in the opening greeting, reviewed against Texas regulations; callers who decline route to a human.
- **Human review** of every call touching an emergency, vulnerable resident, or legal-risk statement within 24 hours.

## Success metrics

METRIC	TARGET	WHY IT MATTERS
Missed calls reduced	≥ 50%	Captures volume that previously went uncaptured
Response time	< 15 sec	Customer-perceived speed
Booking rate	≥ 60%	Dispatcher load reduction
Lead capture	100%	Revenue protection
Manual admin time	≥ 30% less	Team capacity recovery
Escalation accuracy	≥ 90%	Safety and trust

## Go / no-go decision logic

---

### GO

All six KPIs hit target, zero safety incidents, and dispatchers confirm operational value during the weekly sync.

### NO-GO

Any safety incident, escalation accuracy below 80%, or dispatcher feedback that the agent created more work than it removed.

## Appendix - a mock call

---

**AI**

Thank you for calling Enderby Gas. I'm an AI assistant. I can help with deliveries, tank levels, billing questions, and emergency triage, and I'll connect you with a person if you'd prefer. This call is being recorded. How can I help today?

**CALLER**

Yeah, hi. I think we're out of gas. The heat just stopped working.

**AI**

I'm sorry to hear that. Before we go further, is anyone in the home elderly, an infant, or in a medical situation where the cold is an immediate concern?

**CALLER**

No, it's just me and my husband. We're fine, just cold.

**AI**

Thank you. And do you notice any gas smell, hissing sounds, or anything unusual?

**CALLER**

No, nothing like that. The tank's just empty I think.

**AI**

Got it. Can I have the service address please?

**CALLER**

4421 County Road 217, Gainesville.

**AI**

Thank you. I'm sending a same-day delivery request to dispatch now. Someone will follow up by phone within 30 minutes with a delivery window. In the meantime, please don't attempt to relight any pilots or appliances yourself - a technician will handle that when the tank is filled. Anything else?

**CALLER**

No, that's it. Thank you.

**AI**

You're welcome. Stay warm. Goodbye.

**Outcome:** Call summary auto-generated and routed to the dispatcher queue with a priority flag (no-heat, no safety risk, residential). Caller receives SMS confirmation within 60 seconds.

## Why it transfers to startups

This shows the ability to scope and de-risk an AI deployment end-to-end - safety guardrails, measurable success criteria, human-in-the-loop review, and an honest go/no-go framework. That's senior-level product and operations judgment, not just wiring tools together.

---

*Enderby Gas is a real company used with the owner's permission. All names, customer data, and operational figures are simulated for portfolio purposes; cost and impact figures are based on publicly available pricing and industry benchmarks. No proprietary records were used.*