


ONTIQ FIELD GUIDE



# AI for the Front Office

How an AI receptionist, an ontology and an edge box  
fit together — written for owners, not engineers.

---

Ontiq Labs · [ontiqlabs.com](https://ontiqlabs.com) · [info@ontiqlabs.com](mailto:info@ontiqlabs.com)

Worked example inside: Silver Creek Ranch Resort, Hill Country TX



## CHAPTER 1

## The 6:00 PM problem

Every service business has the same hidden leak: the phone. Calls after hours go to voicemail — and voicemail is where revenue goes to die. Calls during the rush stack up while your best person juggles a guest at the desk, a text about a late checkout, and a contractor asking for a gate code. None of this is a people problem. It is an information problem: the answers exist — in your PMS, your booking engine, your maintenance log, your policies — but no human can hold all of it at 6:00 PM on a Friday.

The AI industry's answer has been a chatbot bolted to a website. Ours is different: give the machine the same thing you'd give a great new hire — **a model of how your business actually works** — then let it answer the phone with that model behind it.

## CHAPTER 2

## What an ontology is (in plain English)

An ontology is your business, written down as things and rules. Not documents — **things**: Guests.

Bookings. Units. Trail rides. Waivers. Work orders. And the relationships between them: a Guest *holds* a Booking; a Booking *occupies* a Unit; a trail ride *requires* a signed Waiver.

Once your operation exists in this form, a machine can answer questions the way your ten-year veteran does — by **looking things up and following the rules**, not by guessing. That distinction is the entire product. A language model guesses beautifully. An ontology doesn't guess at all.

**Graph first. Tools second. Tokens last. In production, 92% of front-office questions never touch an AI model — they're answered straight from the graph, at \$0.**

## CHAPTER 3

## The stack, in one page

**The edge box.** A small appliance in your building runs everything: the voice agent, the ontology, local AI models, and the security stack. Your data never leaves. The cloud is our supply line — signed model updates travel in; your guest data doesn't travel out.

**The receptionist.** Answers every call in under a second, speaks from the graph, books, reschedules, texts back — and hands to a human the moment a caller is frustrated or a rule says “this needs a person.”

**The guardrails.** Every answer is checked against the ontology before it's spoken. Every action is logged. Anything the graph can't verify, the agent won't claim.

## WORKED EXAMPLE

## Silver Creek Ranch Resort



*214 units across cabins and lodge rooms · a spa · stables · three dining rooms · Hill Country, Texas. Names changed, numbers real-world typical.*

Silver Creek's front desk handled 300+ calls a day in season. A third were the same twelve questions. Trail rides required signed waivers — chased by hand. After 9 PM, calls went to a binder and a prayer.

### Their ontology, at go-live:

| Entities                           | Links                     | Actions                          | Rules                               |
|------------------------------------|---------------------------|----------------------------------|-------------------------------------|
| Guest · Booking · Unit · Party     | Guest holds Booking       | book · reschedule · cancel       | No trail ride without signed waiver |
| Activity · Waiver · Task · Check   | Booking occupies Unit     | send waiver link by text         | Quiet hours: no outbound after 9 PM |
| Rate · Policy · Staff · Vendor     | Activity requires Waiver  | open + route work orders         | Pets: cabins yes, lodge no — always |
| (grows weekly from real questions) | Tickets assigned-to Staff | escalate to human on frustration | Refunds over \$200: human approves  |

### A Tuesday, 7:02 AM

A guest calls: **“Can my daughter ride this afternoon if my husband signed her waiver last year?”**  
The graph checks: Waiver → signed 14 months ago → policy says annual renewal → agent answers in one breath, texts a fresh waiver link, holds two saddles, and logs it. Zero staff seconds. The model never guessed — it looked it up.

By month two: **96% of calls answered inside one ring · the twelve repeat questions fully absorbed · waiver chase time to zero · an estimated \$11k/month in recovered after-hours bookings.** The ontology grew from 100 covered questions to 109 — because the corpus loop watches what real callers actually ask.

## CHAPTER 5

# Security is the foundation, not a feature

We are a security company that built an AI platform, not the reverse. Every endpoint in your building — desk PCs, door controllers, cameras, the edge box itself — carries enterprise EDR and reports to a 24/7 SOC with 40-second containment playbooks. Your design files, guest lists and call recordings never leave your building. Staff get one governed way to use AI — so they stop pasting your business into public chatbots.

## CHAPTER 6

# The door opens from the inside

Everything above runs on hardware you can own, models you can buy outright, and an ontology that is **yours** — exportable, documented, portable. We feel like you'll never leave us. But if you do: you're welcome. That promise is on our pricing page, in writing, because a platform you can leave is the only kind worth staying on.

**Next step:** a 30-minute walkthrough prices your exact stack — fixed quote, in writing, same day.  
info@ontiqlabs.com · ontiqlabs.com