



SERVICE GUIDE · 2026

# Conversational AI

Answers your customers and your  
staff can actually trust.

## INSIDE THIS GUIDE

- 01 How it works
- 02 What it can actually do
- 03 Where it applies
- 04 What we build for you
- 05 Accuracy & governance
- 06 Hosting & deployment
- 07 Technical companion guide
- 08 Training & support
- 09 Common questions
- 10 Our delivery process

## ENTRA AI

[www.entra-ai.com](http://www.entra-ai.com)  
[info@entra-ai.com](mailto:info@entra-ai.com)  
+44 2034 323178

## WHAT IS CONVERSATIONAL AI

# Answers, grounded in your own truth

Everyone has met a bad chatbot: the one that offers four buttons, none of which is your question, and then apologises. The technology has moved on considerably — but the thing that decides whether an assistant is useful is not how well it writes.

**It is whether the answer is actually right.** We build assistants that answer from *your* documents, *your* policies, and *your* live systems — with a link back to the source, and with the good sense to hand over to a human when they are out of their depth.

## THINK OF IT LIKE THIS

*The difference is between a confident stranger who has read about your industry and a new starter who has read your actual handbook, has your systems open in front of them, and knows to ask a colleague when they are unsure. We build the second one.*

## 01

## How It Works

A plain-English overview

### ■ It Reads Your Material

Your documents, policies, product data, and past cases become the assistant's source of truth — kept current as they change.

### ■ It Answers From Source

Every response is built from what it found in your material, with the reference available so anyone can check it.

### ■ It Knows Its Limits

When the answer is not in your sources or the request needs authority, it hands over to a person with the context attached.

## 02

## What It Can Actually Do

Beyond answering questions

### ■ Customer Self-Service

Genuine answers to real questions at any hour, on your website, in your app, or through the channels customers already use.

### ■ Agent Assist

Your people stay in the conversation, but with answers, history, and a suggested response already prepared.

### ■ Onboarding & Training

New starters and new customers guided through processes conversationally, at their own pace.

### ■ Internal Knowledge Assistant

One place your staff can ask about policy, process, products, or systems instead of hunting through six intranet pages.

### ■ Transactional Requests

Not just answering but doing — checking an order, booking a slot, raising a ticket, updating a record.

### ■ Multilingual Support

The same grounded answers in the languages your customers and your teams actually use.

## 03 · WHERE IT APPLIES

# Real scenarios, real outcomes

Conversational AI pays off wherever the same questions are answered repeatedly by people whose time is worth more than that. Below are the deployments we are asked for most often.

## A

## Customer-Facing Support

Faster answers without a bigger team

### ■ First-Line Support

The high-volume, well-documented questions handled properly, so your team is left with the ones that genuinely need them.

### ■ Product & Service Guidance

Helping customers find the right option, with answers drawn from current specifications rather than last year's brochure.

### ■ Order & Account Enquiries

Live answers pulled from your actual systems — status, history, entitlements — rather than a generic help article.

### ■ Out-of-Hours Cover

Real assistance overnight and at weekends, with clean escalation waiting for your team in the morning.

## B

## Internal & Employee Support

Ending the search through six intranet pages

### ■ HR & Policy Questions

Leave, expenses, benefits, and procedure answered consistently from your current, approved documentation.

### ■ Sales & Bid Enablement

Instant access to specifications, pricing rules, case studies, and previous answers while the customer is still on the call.

### ■ IT & Service Desk

Common issues diagnosed and resolved conversationally, with tickets raised automatically when they are not.

### ■ Operational Procedure

Frontline staff getting the correct process at the moment they need it, rather than the version they remember.

## C

## Regulated & High-Stakes Settings

Where being confidently wrong is expensive

### ■ Source-Cited Responses

Every answer carries its reference, so any response can be verified against the approved document behind it.

### ■ Full Conversation Records

Complete, searchable transcripts of what was asked and answered — for quality review, audit, and dispute resolution.

### ■ Approved-Content Only

The assistant answers strictly from material you have signed off, and says so plainly when a question falls outside it.

### ■ Controlled Escalation

Defined categories that always route to a qualified person, with the conversation and context carried across.

## 04

## What We Build For You

The parts of a working system

### ■ Knowledge Pipeline

The process that brings your documents and data in, keeps them current, and retires content when it is superseded.

### ■ Conversation Design

Tone, boundaries, and behaviour that fit your brand, including how it says 'I do not know' without sounding useless.

### ■ Escalation & Handover

Clean routes to your team, with the full conversation and context passed across so nobody repeats themselves.

### ■ Retrieval & Grounding Layer

The component that finds the right passage for each question — the single biggest factor in whether answers are correct.

### ■ System Integrations

Secure connections to the systems that hold live data, so it can answer about *this* order rather than orders in general.

### ■ Analytics & Quality Review

What is being asked, what is being answered well, and where the gaps in your knowledge base actually are.

## 05

## Accuracy & Governance

Grounded, cited, and reviewable

### ■ Answers From Your Sources

The assistant is constrained to your approved material. It is not improvising from whatever it read on the internet.

### ■ Clear About Its Limits

When the answer is not there, it says so and offers a person. Designed refusal beats confident invention every time.

### ■ Full Audit Trail

Every conversation recorded and searchable, with the sources used for each answer retained alongside it.

### ■ Citations Everywhere

Each answer links to the document and section it came from, so trust is verifiable rather than assumed.

### ■ Quality Monitoring

We sample and score real conversations continuously, so quality is measured rather than assumed after launch.

### ■ Data Privacy & Security

Your content and your conversations stay yours, within your security policies and residency requirements.

## 06

## Hosting & Deployment

Wherever your policies require

### ■ Your Cloud

Deployed into your own AWS, Azure, or Google Cloud environment under your existing controls.

### ■ On-Premise

Fully self-hosted, including open models, where regulation or policy requires it.

### ■ Managed by Us

We host, monitor, and maintain it and deliver a supported service.

## GO DEEPER · TECHNICAL COMPANION GUIDE

**07 · Retrieval-augmented generation — how grounding works**

The reason our assistants answer correctly is **RAG (retrieval-augmented generation)**: before the model writes anything, the system retrieves the relevant passages from your own material and requires the answer to be built from them. Our technical companion guide covers chunking strategy, embedding and vector search, hybrid retrieval, re-ranking, and evaluation — the detail your architects will want.

[www.entra-ai.com/assets/docs/Entra\\_AI\\_RAG\\_Service\\_Guide.pdf](http://www.entra-ai.com/assets/docs/Entra_AI_RAG_Service_Guide.pdf)

**08****Training & Support**

Making sure people keep using it

Adoption is won or lost on the first few answers people get. If those are wrong, they stop asking and rarely come back — so we launch with a narrow scope done well rather than a broad scope done adequately.

**■ Team Onboarding**

How to use it, how to spot a weak answer, and how to feed corrections back into the knowledge base.

**■ Leadership Briefings**

Coverage, deflection rates, escalation patterns, and what the system is genuinely saving.

**■ Content Ownership**

We train your team to maintain the source material, since answer quality follows content quality directly.

**■ Ongoing Support**

Reviews of real conversations, tuning, and expansion into new topics as confidence grows.

## 09

## Common Questions

The things clients ask us first

### ■ Will it make things up?

It answers from your sources and cites them. Where it cannot find an answer, it says so — that behaviour is designed in, not hoped for.

### ■ Will customers know it is AI?

Yes. We recommend being upfront about it and making the route to a human obvious. Concealment costs more trust than it saves.

### ■ Is our content used to train models?

No. Your content answers your questions and nothing else. That is contractual.

### ■ Our documentation is a mess.

That is normal, and part of the work. We identify the gaps and contradictions early — most clients find that valuable on its own.

### ■ What about the hard questions?

They escalate, with full context attached. The aim is to free your experts for those, not to keep customers away from them.

### ■ How long until it is live?

It depends heavily on the state of your content. We assess that first and agree a realistic timeline before any commitment is made.

## 10

## Our Delivery Process

From first conversation to a live assistant

### ■ 01 — Discovery Call

We learn who is asking what, how often, and what a good answer actually looks like.

### ■ 04 — Tested Pilot

It runs on real questions with answers reviewed, until quality clears the agreed bar.

### ■ 02 — Content Review

We assess the material the assistant would answer from and flag the gaps up front.

### ■ 05 — Deploy & Train

We go live on your channels and onboard the teams who own it.

### ■ 03 — Build & Ground

We build the retrieval layer and integrations, and tune the conversation design.

### ■ 06 — Monitor & Improve

We review conversations, close knowledge gaps, and extend coverage.



**NEXT STEP**

# Ready to give everyone a straight answer?

Tell us the questions your team answers over and over. We will show you what a grounded assistant would do with them — and where it should hand over to a person.

## GET IN TOUCH

### WEB

[www.entra-ai.com](http://www.entra-ai.com)

### EMAIL

[info@entra-ai.com](mailto:info@entra-ai.com)

### PHONE

+44 2034 323178