



SERVICE GUIDE · 2026

AI Agents & Workflow Automation

Software that does not just answer questions — it does the work.

ALSO KNOWN AS AGENTIC AI

INSIDE THIS GUIDE

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

ENTRA AI

www.entra-ai.com
info@entra-ai.com
+44 2034 323178

WHAT ARE AI AGENTS

From answering to doing

Most AI you have met so far *responds*. You ask, it replies, and a person still does the actual work. An **AI agent** is the step after that: software that can be given a goal, work out the steps, use your systems to carry them out, and come back when the job is done or when it needs a decision from you.

That is what turns AI from an interesting tool into **capacity**. A process that had someone clicking between five systems becomes something that runs quietly in the background, consistently, and at whatever hour the work actually needs doing.

You may hear this called Agentic AI

Agentic AI is the industry term for what we describe here as AI agents — same technology, same service. We lead with the plainer name.

THINK OF IT LIKE THIS

A chatbot is a colleague who answers your questions. An agent is a colleague who takes the task off your desk, does it properly, and tells you what happened — while sticking to the rules you set about what it is and is not allowed to do on its own.

01

How It Works

A plain-English overview

■ You Set the Goal

You describe the outcome you want and the boundaries it must respect — which systems it can touch, and what always needs a human to approve.

■ It Plans and Acts

The agent breaks the goal into steps, uses your tools and data to carry them out, and adapts when something unexpected comes back.

■ You Stay in Control

Every action is logged and visible. Anything sensitive stops and waits for a person to sign it off.

02

What It Can Actually Do

Real work, taken off real desks

■ Multi-Step Processes

Jobs that span several systems and several decisions — the ones too fiddly for traditional automation and too dull for your team.

■ Research & Summarising

Pulling together information from documents, systems, and the web into a briefing that would have taken someone half a day.

■ Triage & Routing

Incoming tickets, emails, claims, or applications read, classified, enriched, and sent to the right place with the right context attached.

■ Data Entry & Reconciliation

Moving and checking information between systems that were never designed to talk to each other, and flagging what does not match.

■ Drafting & Preparation

First drafts of reports, responses, quotes, and summaries — prepared for a person to review and approve rather than write from scratch.

■ Monitoring & Response

Watching for conditions that matter and acting on them — escalating, notifying, or kicking off the next process automatically.

03 · WHERE IT APPLIES

Real scenarios, real outcomes

Below are the patterns we see most often. Each one replaces hours of repetitive human effort, or makes something possible that simply was not worth a person's time before.

A

Operations & Back Office

The processes that quietly consume your week

■ Invoice & Document Processing

Documents read, key details extracted, checked against your records, and pushed into the finance system — with exceptions flagged rather than guessed at.

■ Supplier & Data Onboarding

New records created, validated, enriched from external sources, and checked for duplicates before anything reaches your master data.

■ Order & Case Management

Cases progressed through their stages automatically, chasing missing information and updating everyone who needs to know.

■ Compliance Checks

Routine checks run consistently on every record rather than on a sample, with a full trail of what was checked and when.

B

Customer & Service Teams

Faster responses without more headcount

■ Ticket Triage

Every incoming request read, categorised, prioritised, and routed — with the relevant history and account context already attached.

■ Proactive Follow-Up

Chasing outstanding information, confirming completed work, and spotting cases that have gone quiet before they become complaints.

■ First-Line Resolution

Common requests handled end to end, from the answer itself through to updating the record and closing the ticket.

■ Agent Assist

Your people stay in charge, but with the research, history, and draft response already prepared when they open the case.

C

Knowledge & Decision Support

Work that used to need a specialist's morning

■ Research Briefings

Structured summaries pulled from internal documents and external sources, with every claim traceable back to where it came from.

■ Tender & Bid Support

Requirements extracted, matched against your existing material, and turned into a first draft response for your team to sharpen.

■ Reporting & Commentary

Recurring reports assembled and written up, so your analysts spend their time on the interpretation rather than the assembly.

■ Internal Q&A

Questions about policy, process, or product answered from your own approved sources, with the action taken as well as the answer given.

04

What We Build For You

The parts of a working system

■ Agent Design & Orchestration

The architecture itself — how work is broken down, which agent handles what, and how they hand off between each other.

■ Guardrails & Approval Gates

Explicit rules on what can happen automatically and what always stops for a human decision.

■ Monitoring & Activity Logs

A clear, searchable record of every action, decision, and outcome — for your operations team and your auditors alike.

■ Tool & System Integration

Secure, reliable connections into the systems the agent needs, so it works with your stack rather than beside it.

■ Memory & Context Handling

So the agent knows what has already happened on a case and does not start from zero every time.

■ Evaluation Harness

Automated testing against real historical cases, so you can measure quality before go-live and catch drift afterwards.

05

Oversight & Control

Autonomy with a handbrake

■ Full Visibility

See exactly what the system is doing — every action, decision, and outcome — through a clear activity log.

■ Audit Trails

Every action recorded with a timestamp and a reason: a searchable record for compliance and review.

■ Guardrails & Boundaries

Hard limits on what the agent can and cannot do. Sensitive actions always require human sign-off.

■ Graceful Escalation

When the agent is not confident, it stops and asks rather than guessing. Uncertainty becomes a handover, not an error.

■ Ongoing Monitoring

We track performance over time, flagging anything unusual and confirming accuracy holds as your processes change.

■ Data Privacy & Security

Your data stays yours. We work within your security policies and your data 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 and billing.

■ On-Premise

Fully self-hosted where regulation, data residency, or policy demands it. Nothing leaves your network.

■ Managed by Us

We host, monitor, and maintain it, and you get a working service with an agreed support arrangement.

[GO DEEPER · TECHNICAL COMPANION GUIDE](#)

07 · Model Context Protocol — the plumbing underneath

Agents are only as useful as the systems they can reach. We build those connections on **Model Context Protocol (MCP)**, an open standard for linking models to tools and data. It replaces brittle one-off integrations with a consistent interface, which means less maintenance and far less pain when you change model or vendor later. Our MCP guide is the technical read for your architects and engineers.

www.entra-ai.com/assets/docs/Entra_AI_MCP_Service_Guide.pdf

08

Training & Support

Making sure your team trusts it

An agent that people do not trust gets switched off, however well it performs. We spend real time on the handover so your team knows what it does, when to step in, and how to tell whether it is working.

■ Team Onboarding

Practical, jargon-free sessions for the people working alongside the system every day — what it does, and when to intervene.

■ Leadership Briefings

Focused sessions covering outcomes, governance, risk, and how success is being measured.

■ Runbooks & Playbooks

Plain-English documentation for your system: what to check, what to do when something looks wrong, and who to call.

■ Ongoing Support

As your processes evolve, we are available for adjustments, new workflows, and periodic reviews.

09

Common Questions

The things clients ask us first

■ Will it act without asking us?

Only where you have said it can. Every deployment starts conservative and widens as you build confidence in it.

■ Do we need to replace our systems?

No. Agents work through the systems you already have, using the same interfaces and permissions your people do.

■ Is our data used to train models?

No. Your data is used to do your work, and nothing else. That is contractual, not just policy.

■ What happens when it gets it wrong?

It escalates rather than guesses, and every action is logged and reversible. We design for the bad day, not just the good one.

■ How do we know it is working?

We agree the measures before we build — volume handled, accuracy, escalation rate, time saved — and report against them.

■ How long until it is live?

It depends on the workflow and the systems involved. We scope it properly up front and agree a realistic timeline before we start.

10

Our Delivery Process

From first conversation to a live system

■ 01 — Discovery Call

We learn about the process, the volumes, and where the pain actually sits.

■ 04 — Supervised Pilot

It runs alongside your team with every action reviewed, until the numbers earn your trust.

■ 02 — Scoping & Proposal

We map the workflow, confirm feasibility, and agree the outcomes and the guardrails.

■ 05 — Deploy & Train

We go live on your chosen infrastructure and onboard the people who will work with it.

■ 03 — Build & Test

We build the system and test it against your real historical cases, not toy examples.

■ 06 — Monitor & Improve

We track performance, maintain oversight, and extend the system to the next workflow.



NEXT STEP

Ready to hand the busywork to something else?

Tell us about a process that eats your team's week. We will tell you straight whether an agent should run it, and what it would take to build.

GET IN TOUCH

WEB

www.entra-ai.com

EMAIL

info@entra-ai.com

PHONE

+44 2034 323178