



Your institution's own AI. *Built for your documents and your rules.*



GOVERNABLE SPECIFIC-PURPOSE AI • THE AWEN PLATFORM

We design and build AI systems, one for each specific purpose inside your organisation. Each one answers from your documents, works under the rules you set, and keeps a record you can show to anyone who asks — able to account for every answer.

Your documents

It answers from your files, not from the internet.

Your rules

It does what you allow and nothing else.

Your record

Every answer leaves evidence you can show.

— BEFORE THE TECHNOLOGY

AI answers. It rarely shows its work.

You ask it something about a document. The answer comes back fluent, confident, well written. It may well be correct. Then someone asks you: where did this figure come from? And there is nothing to show.

Awentir returns the answer, the exact page it came from, and a written record of what happened — kept, and yours to show.



◆ BUILT FOR ONE JOB — YOUR DOCUMENTS, YOUR RULES

— HOW WE ENGAGE

One discipline. Three ways to put it to work.

01 · BUILD

Bespoke specific-purpose AI

Your institution's own AI, engineered on the Awen platform for your documents, your rules and your deployment — in your cloud or ours. You receive a working system tuned to your domain and tested on your own cases, the rules it operates under written down and agreed with you, an evidence record for every answer, training and support. Typical projects: 4 to 12 weeks.

◆ TAKING PROJECTS NOW

02 · ADVISE

Advisory

Consulting for large AI programmes: governance-by-design architectures, explainability (XAI) audits of existing systems, regulatory alignment with the EU AI Act, NIST AI RMF and ISO/IEC 42001.

◆ BY ENGAGEMENT

03 · CONNECT

Governance Platform

Bring your own AI. Connect it to Awentir and obtain Defensibility Records — an audit-ready record of what your AI did, on what evidence, under which policy — for the systems you already run.

◆ DESIGN-PARTNER PROGRAMME

Any organisation whose work runs on documents, rules and decisions that may one day be questioned: by an auditor, a regulator, a court or a board.

SECTORS

Where we work

- ◆ Government and public sector
- ◆ Municipalities
- ◆ Health and hospitals
- ◆ Insurance
- ◆ Workplace risk management
- ◆ Banking and finance
- ◆ Legal
- ◆ Public procurement
- ◆ Telecommunications and infrastructure
- ◆ Hospitality
- ◆ Utilities and energy
- ◆ Industry and logistics
- ◆ Education

FUNCTIONS

What we build for

- ◆ Administration, finance and control
- ◆ Internal control and auditing
- ◆ Risk management
- ◆ Compliance
- ◆ Purchasing and planning
- ◆ Legal and contracts
- ◆ Cybersecurity and cyber defence
- ◆ Governable AI for cybersecurity and cyber defence

Your sector is not listed? Talk to us about your documents.

METHOD

From document to defensible claim.

0 1

Set the mandate

You define what the system may use: which documents, which tools, which data zone. The mandate stays visible for the whole session.

0 2

Ground the answer

The system retrieves the relevant passages from your documents — it answers from your evidence, not from memory.

0 3

Claims, with citations

Findings are returned as discrete claims, each bound to its source page or clause. Gold marks evidence — nothing else.

0 4

Keep the record

Which documents were used, which rule applied and which evidence supported the answer are written to the record. What happened is what is recorded.

◆ GOVERNANCE, INCLUDED — EVERY ANSWER ON THE RECORD

GOVERNANCE

Governance, included.

Every system we build comes with its evidence record: what it did, on what evidence, under which rules. That record is yours, exportable, and readable by people who never used the system.

◆ If your regulator asks tomorrow

Under the EU AI Act, ISO/IEC 42001, NIST AI RMF, or a national rule that does not exist yet — the record already exists.

◆ If nobody asks

It is simply good practice: your people know why the system said what it said.

◆ Rules you set, written down

What the system may use and what it must never do is agreed with you before it is built, and stays with the system.

◆ Built for the auditor's question

"Show me why the system said this" has a literal answer: the claim, its citation, the rule applied and the record — exportable as evidence.

What does not change.

- ◆ Your documents stay in your data zone and are never used to train models.
- ◆ Accepting a finding into your work remains a human act, by design.
- ◆ Gold marks evidence. Nothing else.

Research-led. The foundations of specific-purpose AI.

RESEARCH LINE · 01	Explainable AI (XAI) Turning model explanations into evidence a professional can defend.
RESEARCH LINE · 02	Specific-purpose AI Domain-bound systems that work under a mandate and answer from documents, not from memory.
RESEARCH LINE · 03	Agent orchestration Coordinating several AI agents on one task with human oversight and a record of who did what.
RESEARCH LINE · 04	Governance of AI systems Evidence, records and controls aligned with the EU AI Act, NIST AI RMF and ISO/IEC 42001.
RESEARCH LINE · 05	Governable AI for smart hospitals Clinical and administrative systems that answer from records and protocols, with human oversight and a full record.
RESEARCH LINE · 06	AI-enabled data collection and correlation for cyber defence Governable pipelines that gather, correlate and explain security signals, keeping every conclusion traceable to its evidence.

◆ GOVERNANCE, INCLUDED — EVERY ANSWER ON THE RECORD

Start a conversation.

Tell us what your organisation does, which documents it runs on and what you need the system to do. We reply within two working days, and the first conversation is free.

WRITE TO US

info@awentir.ai

AWENTIR.AI



SCAN TO SAVE OUR CONTACT

AWENTIR — A governable specific-purpose AI company: services, bespoke, advisory and a governance platform — on the Awen platform. **Gold marks evidence. Nothing else.**

AWENTIR LTD · Registered in England and Wales · Company number 17423584 · Registered office: 20 Wenlock Road, London N1 7GU, United Kingdom · info@awentir.ai · © 2026 Awentir