

An internal AI capability without hiring an AI department

Ongoing support for SMEs that want to turn processes, data and company knowledge into agents, RAG systems, automations and faster decisions.

Analysis -> Roadmap -> Development -> Training -> Governance



The problem: AI is purchased, but rarely integrated.

Many SMEs start with tools and demos, then get stuck because there is no technical reading of their real processes.

No roadmap

It is not clear what to automate, with which priorities and which results to measure.

Fragmented systems

CRM, ERP, sheets, email, documents and databases do not communicate naturally.

Sensitive data

Privacy, NDAs, roles and permissions require the right architecture choices.

Uncontrolled costs

APIs, tokens, cloud, storage, servers and maintenance must be planned in advance.

Weak adoption

Without training and control, the team does not use AI consistently.



The proposal: an AI Architect in Residence.

An external but ongoing role that enters the company's working method and builds by priority.

AIR does not sell random on-demand hours. It brings a sequence: diagnosis, roadmap, progressive development, training and results control.

- 01 Reviews processes, software, data and constraints.
- 02 Defines what to build first and what to avoid.
- 03 Builds agents, RAG systems, automations and dashboards.
- 04 Trains the team and documents operating rules.



The method reduces risk before building.

We start with a paid assessment, useful even on its own. If the client activates AIR, the cost can be deducted from the contract.

01

AI Readiness Check

Processes, software, data, people, constraints, priorities.

02

Agentic roadmap

Agents, RAG, automations, governance, KPIs and risks.

03

Monthly support

8-40 hours/month to design, develop and train.

04

Continuous evolution

Measurement, feedback, new LLM versions, new needs.

Guided decision, not a leap in the dark

The client buys clarity first. Then, if it makes sense, they buy operational continuity.

AIR starts from real work, not from technology.

First we understand where time is lost, where errors arise and which information has decision-making value.

Processes and departments

Sales, administration, operations, marketing, customer care, management.

Software and integrations

CRM, ERP, business software, Google Workspace, Microsoft 365, APIs and databases.

Data and documents

PDFs, contracts, manuals, emails, sheets, reports, procedures and archives.

Governance and privacy

Roles, permissions, NDAs, logs, audits, supervision and internal policies.

Costs and sustainability

Tokens, cloud, local models, servers, licenses and maintenance.



What gets built: enterprise agentic infrastructure.

Not generic chatbots, but systems connected to data, procedures and operational tools.

Company sources

Documents, emails, sheets, CRM, ERP, databases, procedures.

AI layer

RAG, vertical agents, prompt library, workflows, guardrails.

Operational outputs

Reports, dashboards, alerts, automations, controlled responses.

Governance

Permissions, logs, audits, human control, data policies.



The commercial difference: every asset built remains useful to the company and can evolve over time.

The arguments that convince a B2B client.

AIR answers the most common objections: cost, control, privacy, adoption and results.

Predictable budget

Monthly plans with clear hours: the client knows the investment before starting.

Progressive ROI

Work is prioritized: first the processes with the clearest impact on time, errors or speed.

No internal AI department

The company gains specialist expertise without immediately hiring expensive roles that are hard to coordinate.

Privacy and control

Cloud, on-premise or hybrid architecture are chosen based on data, risk and budget.

Real adoption

Training, procedures and documentation prevent AI from remaining an unused demo.

Why AIR is better suited to SMEs than the alternatives.

It is an intermediate model: more stable than spot consulting, less costly than an internal department.

Criterion	Spot consulting	Internal hiring	Ergoware AIR
Startup speed	High, but often isolated	Low: recruiting and onboarding	High, with method and initial diagnosis
Fixed cost	Limited, but discontinuous	High and permanent	Monthly, modular, predictable
Process knowledge	Partial	High over time	Grows month after month
Delivery tecnica	Not always included	Depends on the hired team	Strategy + development + training
Governance AI	Often unstructured	To be built internally	Part of the service

Commercial positioning: AIR does not replace management. It provides technical expertise, method and execution continuity.

Clear plans: from light support to near-internal presence.

Reference hourly rate €90/h, presented as a monthly plan to make the budget predictable.

Essential

8 hours/month

€720

Analysis, priorities, small prototypes, targeted training.

Growth

16 hours/month

€1,440

Roadmap, simple agents, RAG and progressive automations.

Advanced

24 hours/month

€2,160

More departments, more data, more workflows to transform.

Embedded

40 hours/month

€3,600

Near-internal support, continuous development and governance.

Not automatically included: licenses, APIs, tokens, cloud, hardware, servers, storage, travel, unplanned system maintenance and extra development beyond the monthly hours.

First step: understand how many hours are really needed.

The client does not have to choose blindly. We start from a concrete assessment and choose the plan only after reading the context.

1

Exploratory call

industry, size, tools, critical areas

2

AI Readiness Check

processes, data, constraints, opportunities, risks

3

Operational roadmap

priorities, plans, hours, expected outputs

4

AIR activation

progressive development, training, governance



Request an assessment

ergowareltd.com/ai-architect-in-residence.html