

# GARION

Organizational AI Readiness Intelligence

*AI readiness. Governed. Proven.*

---

Whitepaper

## **Advisors & the EU AI Act**

*From conversation to a robust mandate: how advisors, coaches and consultants build a scalable offering around AI readiness with GARION.*

GARION Consultant Edition · Local macOS and Windows application · German & English  
As of June 2026 · prega-design · predrag.gasic@prega-design.de

# Contents

1. Management summary
  2. The market — why a consulting window is open now
  3. The problem GARION solves for advisors
  4. What GARION is
  5. The methodology — your professional backbone in the pitch
  6. The Consultant Edition
  7. Your business model — the fee ladder
  8. What you actually deliver
  9. Trust, delimitation and data protection
  10. Objection handling
  11. Next steps
- Legal notice & sources

A note on terminology: „EU AI Act“ is the common name for Regulation (EU) 2024/1689 (in German: „EU-KI-Verordnung“). Both terms refer to the same law.

# 1. Management summary

## In 2026 the EU AI Act is no longer a question for the future, but operational reality.

The AI literacy obligation under Article 4 has applied since 2 February 2025; from 2 August 2026 the national market-surveillance authorities receive the formal powers to enforce it. For every organization that uses or provides AI, a concrete, evidence-bound need for action arises — and for you as an advisor, a precisely timed market.

Your clients' real problem is rarely the will, but the proof: who can demonstrate today that the right roles hold the right AI literacy, that risks are documented, and that management can defend this before an auditor? This is exactly where GARION starts — and exactly where your fee lies.

This whitepaper is addressed to advisors, coaches and consultants. It shows the market mechanism, the methodology behind GARION and, above all, the business model with which you turn it into a viable offering.

### The three questions that decide a mandate

Management teams do not want an 80-page opinion, but a clear answer to three questions: Where do we stand — in one number? Where is our biggest risk? What is the next step — and what does it achieve? Whoever answers these convincingly in a first meeting wins the mandate.

## 2. The market — why a consulting window is open now

### 2.1 The obligation is already active

Article 4 of the EU AI Act obliges providers and deployers of AI systems to ensure, „to their best extent“, a sufficient level of AI literacy among their staff and all persons who handle AI systems on their behalf. Unlike the high-risk requirements, this obligation applies to practically every company that uses AI — not only the Annex III cases.

Note: in the course of the „Digital Omnibus“, the wording of Article 4 is expected to be softened — from „ensure“ towards „support the development of AI literacy“. For deployers of high-risk systems, the training obligation for human oversight remains. The message to clients does not change: whoever cannot demonstrate competence and documentation carries an avoidable risk.

### 2.2 The timeline a client needs to hear

Date	What happens
2 February 2025	The AI literacy obligation (Art. 4) and the prohibitions on unacceptable practices apply.
2 August 2026	Market-surveillance authorities receive the formal enforcement powers.
from 2027	Further high-risk obligations take effect; the „Digital Omnibus“ may adjust deadlines.

Regulatory deadlines create consulting demand — but only if the offer is concrete, fast and presentable. Whoever shows up with a concrete, presentable instrument instead of a 90-day audit wins the mandates. The others are still explaining the problem while you are already delivering the solution.

### 3. The problem GARION solves for advisors

Most advisors in the AI and governance space run into the same four bottlenecks — they limit growth and margin:

- **Long ramp-up.** Data collection and preparation per client take weeks — this caps the number of mandates.
- **Quality depends on the person.** Results vary with the day and seniority; this makes scaling harder.
- **No presentable result.** Without a board-ready artifact, the first meeting leads nowhere.
- **No recurring revenue.** Project spikes instead of predictable income.

GARION resolves all four at once: a result in the first session, deterministic and therefore consistent quality, a premium report to hand over, and a model for quarterly re-assessment.

### 4. What GARION is

GARION is a local macOS and Windows application that turns scattered facts about an organization, people, roles and evidence into a single, explainable and audit-ready picture of AI readiness — and tells leadership what to do next. Four nouns are enough to understand the model:

Concept	Meaning
Organization	The entity being assessed (one per database).
Role	A function with an AI usage level, a risk level and a required level of competence.
Evidence	A proof of readiness: a training, a certification, a policy acceptance.
Readiness	The computed verdict: how well evidence covers the roles, weighted by risk.

#### 4.1 Four properties that count in the sales conversation

- **Local & confidential.** No cloud, no server. All data lives in a local file. The only optional network call is an explicitly approved AI explanation.
- **Deterministic.** The same inputs always yield the same number. The score is never „invented“ by the AI — the AI only formulates the explanation. This makes every statement defensible.
- **Audit-ready.** Every figure is traceable to the underlying evidence; every change and every AI call is logged.
- **Bilingual.** Interface, reports and AI texts in German and English — natively generated, not translated.

## 5. The methodology — your professional backbone in the pitch

**GARION follows a clear doctrine: Readiness → Risks → Gaps → Recommendations → Details. This is at the same time the dramaturgy with which you can run a first meeting.**

### 5.1 The readiness score

A single value from 0–100 with a named band, plus a confidence level and a trend versus the last measurement. High-risk roles without competence evidence lower the score about three times more than low risks — reflecting the risk-based spirit of the EU AI Act without claiming conformity.

Score	Band	Meaning
0–39	At Risk	Urgent need for action.
40–59	Building	Foundations present, gaps open.
60–79	Developing	Solid, to be sharpened in a targeted way.
80–100	Ready	Robustly documented.

### 5.2 Gaps and recommendations

A deterministic rules engine produces a prioritized list of gaps — such as „high-risk role below the required competence level“ — each with a concrete, pre-written recommended action and an estimated score uplift. The diagnosis immediately becomes an action plan that you, as the advisor, sell and accompany.

### 5.3 The EU AI Act Readiness Check

GARION contains a purely deterministic readiness check across ten areas (incl. AI literacy documented, role-based AI usage, required evidence assigned, policy acceptance, responsibility/owner, risk level, identified gaps, recommended actions, audit trail, non-legal notice). It delivers a status per area — Strong · Partial · Weak · Not enough data.

This clean delimitation is your protection and your selling point at once: you deliver a robust readiness signal without making a conformity promise you cannot keep.

## 6. The Consultant Edition

GARION comes in two editions with an identical core. The Business Edition assesses a single organization. The Consultant Edition wraps a client workspace around the same core — for advisors who serve many clients.

Crucially, the Consultant Edition is a shell, not a re-architecture. Every client is physically isolated — its own database, its own files, no cross-client access. This is exactly the confidentiality your clients demand, and it can be handed over as a portable package when a mandate ends.

## 7. Your business model — the fee ladder

GARION is not a cost factor but a revenue engine. It enables a productized service ladder — from a low-threshold entry to a recurring fee.

Step	Service	Revenue model
Readiness check	Fixed-price initial assessment with Executive One-Pager.	One-off, low threshold.
Remediation mandate	Closing the prioritized gaps (training, owners, policies).	Project, well delimited.
Quarterly retainer	Quarterly re-assessment; trend and current proof.	Recurring, predictable.

### 7.1 Why the margin rises

- **More mandates per advisor.** The ramp-up time per client drops from weeks to minutes.
- **Consistent quality.** The deterministic methodology decouples quality from seniority: junior delivery, senior pricing.
- **White-label effect.** Premium reports raise your perceived value and justify your fee.
- **Predictable revenue.** Retainers based on quarterly re-assessment create recurring income instead of project spikes.

Recommended core offer (advisors): GARION Professional — app + 1 day of consulting, EUR 3,900 net. Entry from EUR 1,900; Premium and Partner packages with positioning, sales story and support up to EUR 9,800. All prices net, plus VAT, as of 2026.

## 8. What you actually deliver

GARION produces a range of deterministically assembled, audit-ready results — each with a provenance and disclaimer block:

- **Executive One-Pager** — a single-page, board-ready PDF with score, confidence, top-3 risks/gaps/actions and EU AI Act Readiness status.
- **AI-Readiness Report** — the complete, explainable scorecard.
- **Audit-Readiness Report** — „Could we export a defensible report now?“ with blocking items and one-click fixes.
- **Role reports** — detailed per function (e.g. marketing, sales, HR, management).
- **Export formats** — PDF for leadership, XLSX matrices for detail work, CSV/Markdown for further processing; .garion backup for handover.

Every report is in German or English — and works fully even with AI switched off (deterministic text building blocks instead of an AI narrative).

## 9. Trust, delimitation and data protection

Especially in a regulatory environment, credibility is everything. GARION is built so that you never promise more than you can keep:

- **No legal advice, no certification.** Every report makes clear that GARION delivers an advisory readiness assessment and does not certify conformity.
- **Human confirms, machine computes.** Scores are deterministic; AI is additive and removable. The final interpretation is confirmed by a human — you.
- **Local & isolated.** No cloud, one database per client. Import runs entirely locally.
- **Full traceability.** Append-only audit log for every relevant change and every AI call.

## 10. Objection handling

Objection	Answer
„Is this legal compliance?“	No — and that is the point. GARION delivers an advisory readiness signal and the evidence trail behind it. The separation from a compliance promise protects you and the client.
„Does the AI invent the score?“	Never. All core values are deterministic; the AI only formulates the explanation and can be switched off.
„Where does our data go?“	Nowhere. Local-first, one database file on the device, no cloud, no account.
„We don't have time.“	You don't need it. A meaningful initial assessment takes under 20 minutes.
„What if the rule changes?“	The obligation to prove competence and documentation remains. GARION builds exactly this evidence and shows a trend.

## 11. Next steps

1. **Choose a pilot mandate.** A manageable organization that already uses AI — ideal for a first assessment in under 20 minutes.
2. **Run the first meeting as a readiness check.** Score, biggest risk, next step — live, closing with an Executive One-Pager.
3. **Offer the fee ladder.** From the fixed-price check to the quarterly retainer.
4. **Build a client base.** Each client in its own isolated workspace — scalable without mixing.

### Why now is the right time

The 2 August 2026 deadline is approaching. Organizations that keep no proof today are looking for guidance now. Whoever establishes a productized readiness offer early wins the mandates — and builds the client base that will need help again at the follow-on deadlines in 2027.

Contact: [predrag.gasic@prega-design.de](mailto:predrag.gasic@prega-design.de) · prega-design

## Legal notice & sources

GARION delivers an advisory assessment of organizational AI readiness. It is not legal advice and not a certification of conformity with the EU AI Act or any other regulation. The Readiness Check assesses the quality of evidence and documentation, not legal conformity. All score values are deterministic; AI-generated texts are marked as advisory and do not replace a professional or legal review. The final interpretation is to be confirmed by a human. Statements on the EU AI Act reflect the status of June 2026 and may change due to ongoing legislative processes (incl. the „Digital Omnibus“).

- EU AI Act — Article 4 (AI literacy): <https://artificialintelligenceact.eu/article/4/>
- AI Act Service Desk (European Commission) — Article 4: <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-4>
- European Commission — AI Literacy, Questions & Answers: <https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers>
- Gibson Dunn — EU AI Act Omnibus Agreement: <https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/>

**ADDENDUM · JUNE 2026**

## **New in GARION: NIS2 governance visibility and controlled data import**

GARION remains what it is: a system for organizational AI readiness intelligence. Since this whitepaper was first published, the platform has grown in two areas, described here factually: an NIS2 governance module for cyber risk visibility and a controlled import of existing organizational data. Only implemented capabilities are described — no roadmap, no promises.

### **NIS2 governance visibility**

Boards and security leaders need an answer to “What should I care about right now?” — in seconds, not after weeks of spreadsheet consolidation. The NIS2 module connects cyber risk data into one readable governance chain:

#### **Risk → Treatment → Evidence → Deadline → Executive summary**

- Risk register: structured cyber risks with title, impact, likelihood, owner and lifecycle status — stored locally in your workspace.
- Treatments: each risk links to one or more responses with owner, status and decision traceability.
- Evidence binding: treatments bind to real evidence records (primary, required, supporting) — proof connected to action.
- Deadline governance: treatments carry obligation windows with unambiguous states — pending, at risk, met, breached — evaluated deterministically in UTC. A met or breached deadline is a fact, not a matter of interpretation.
- Executive summary: a governance card with counts only — risks, open risks, treatments, evidence bindings, deadline states. Orientation in under ten seconds, bilingual (EN/DE).

Deliberate boundaries — by design, not omission: the NIS2 module produces no scores, no charts and no AI interpretation. It is not legal advice, not NIS2 certification and not incident management. GARION shows where you stand — how you act is your decision.

### **Controlled data import**

Nobody wants to type an organization in by hand. GARION imports existing data via CSV — roles, evidence, departments, workforce identities, role assignments and training records — through one controlled pipeline:

- Preview before write: every file is validated and fully displayed before anything is stored. Only explicit approval commits data — no silent writes, no auto-import.
- Deterministic validation: missing required fields, duplicates, unknown references and invalid status values block the commit until resolved. The approval button is physically disabled while blocking errors exist.
- Pseudonymous by architecture: workforce data imports without clear names — opaque identifiers only. Forbidden columns such as first name, birth date or salary reject the whole file before anything is stored. Data minimisation is schema here, not policy.
- Audit and retention: every import is immutably logged with operator, file, counts and outcome; staged data is purged after a successful commit and expires after seven days at the latest.

- Demo and productive cleanly separated: on first launch the organization explicitly chooses between a demo workspace (sample data “Aurora Industries”) and a productive workspace — productive workspaces start empty, with no demo data and no cleanup burden.

### **What this means for your engagements**

For advisors the conversation widens: alongside AI readiness under the EU AI Act you can make NIS2-style cyber risk governance visible — with the same calm, evidence-based methodology. Data import lowers the entry barrier at the client considerably: existing HR, LMS or compliance exports flow into the workspace in a controlled way instead of being typed up in workshops. The demo workspace remains your pitch tool; client data lives separately in productive Consultant Edition workspaces.

All capabilities described are local, license-protected and bilingual (English/German). Complete capability and limitation sheets are available on request. Contact: [office@garion-ai.de](mailto:office@garion-ai.de)