# GARION

Organizational AI Readiness Intelligence

*Understand, prove and document AI readiness.*

---

Whitepaper for companies

# The EU AI Act in Practice

*A factual guide: obligations under the EU AI Act, what „sufficient AI literacy“ means, and how an organization documents its AI readiness in a verifiable way.*

Local macOS and Windows application · German & English · As of June 2026

prega-design · predrag.gasic@prega-design.de

# Contents



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. This whitepaper is for information and does not constitute legal advice.

# 1. About this whitepaper

This whitepaper is addressed to organizations that use or provide Artificial Intelligence and want to meet their obligations under the EU AI Act. It describes factually which obligations arise for deployers, what „sufficient AI literacy“ means in practice, and how an organization records, assesses and documents its position in a verifiable way.

The second part shows how the GARION application supports this work operationally — as a local tool that turns roles and evidence into an explainable readiness picture, prioritizes gaps and produces audit-ready reports. GARION makes no legal statements and does not replace a legal or professional review; it creates the structured factual basis on which such a review can build.

Adressees of this whitepaper are in particular executive management, AI and compliance officers, data protection and IT functions as well as HR and learning & development. It is structured so that chapters 2 to 6 explain the regulatory starting point, chapters 7 to 15 describe how the tool works, and chapters 16 to 20 guide practical implementation. Those who want to start quickly will find the recommended workflow in chapter 16 and a compact checklist in chapter 20.

The description of the regulation reflects the status of June 2026. Ongoing legislative processes (in particular the „Digital Omnibus“) may change individual wordings and deadlines; the fundamental obligation to document AI literacy remains unaffected.

# 2. The EU AI Act at a glance

The EU AI Act is the first comprehensive set of rules for Artificial Intelligence. It follows a risk-based approach: the more an AI system can affect the rights, safety or health of people, the higher the requirements. The regulation distinguishes four risk classes:

| Risk class | Meaning (overview) |
|---|---|
| Unacceptable risk | Prohibited practices (e.g. certain forms of manipulation or social scoring). |
| High risk | Annex III systems with strict requirements for risk management, data, oversight and documentation. |
| Limited risk | Transparency obligations (e.g. labelling of AI-generated content or chatbots). |
| Minimal risk | The majority of applications; no specific obligations beyond the cross-cutting rules. |

The regulation also distinguishes roles: a „provider“ develops or places an AI system on the market; a „deployer“ uses an AI system under its own responsibility. Most companies are in the role of the deployer — and it is precisely for this role that the AI literacy obligation is particularly relevant, because it applies regardless of the risk class.

> **Why Article 4 applies to almost everyone**
>
> Key distinction for practice: the high-risk requirements affect a clearly delimited set of systems. The AI literacy obligation (Article 4), by contrast, affects almost every organization that uses AI at all.

## 2.1 Obligations for deployers of high-risk systems

If an organization uses a high-risk system under Annex III (for example in the areas of employment and recruitment, access to essential services, education or critical infrastructure), additional obligations apply to it as a deployer. These are organizational in nature and overlap strongly with what GARION documents in a structured way:

| Obligation (deployer) | What it is about |
|---|---|
| Human oversight | Assign oversight to suitable persons with the necessary competence and authority. |
| Use according to instructions | Operate the system in line with the provider's instructions for use. |
| Monitoring & reporting | Monitor operation and report incidents or risks. |
| Retention of logs | Keep automatically generated logs for the intended period. |
| Information of affected persons | Inform affected persons of the use where applicable. |

The training and competence requirement for human oversight remains in place even after the planned changes of the „Digital Omnibus“. Whoever keeps the relevant roles, their risk and their competence evidence in good order at the same time fulfils a substantial part of these deployer obligations.

## 2.2 Transparency obligations at limited risk

Independent of the high-risk classification, transparency obligations apply: AI systems that interact with people (such as chatbots) must be recognizable as such; AI-generated or manipulated content (including „deepfakes“) must in principle be labelled. For organizational readiness this means that roles involved in content creation or customer interaction should also be recorded and backed with policies.

## 2.3 General-purpose AI models (GPAI)

For general-purpose models — such as large language models — there are dedicated obligations that primarily affect the providers of these models (technical documentation, copyright, transparency). For most companies they are indirectly relevant: they use such models as deployers and should document in which roles they are used and what competence is required for that.

# 3. Article 4 — the AI literacy obligation

Article 4 obliges providers and deployers to take measures to ensure, to their best extent, a sufficient level of AI literacy among their staff and other persons dealing with the operation and use of AI systems on their behalf. In doing so, the technical knowledge, experience, education and training of the persons as well as the context in which the systems are used must be taken into account.

In practice this means two things. First, the people who work with AI must understand it sufficiently to use it responsibly and appropriately in their respective context. Second, the organization must be able to demonstrate that it made the effort — „to their best extent“ is not „nothing at all“, but proportionate, role-appropriate and documented effort.

## 3.1 Note on the „Digital Omnibus“

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. For operational preparation little changes: an organization that demonstrates competence and documentation in a structured way is in any case better positioned than one that does not.

# 4. Deadlines and timeline

| Date | What happens |
|---|---|
| 2 February 2025 | The AI literacy obligation (Art. 4) and the prohibitions on unacceptable practices become applicable. |
| 2 August 2025 | Obligations for general-purpose AI models (GPAI) and governance structures take effect. |
| 2 August 2026 | National market-surveillance authorities receive the formal enforcement powers; further obligations become applicable. |
| from 2027 | Further requirements for high-risk systems take effect in stages; the „Digital Omnibus“ may adjust individual deadlines. |

For most companies, 2 August 2026 is the decisive reference point: from this date supervisory authorities can enforce the AI literacy obligation. Whoever keeps no robust documentation by then carries an avoidable risk — not because competence is missing, but because it cannot be proven.

# 5. What „sufficient AI literacy“ means in practice

„Sufficient“ is not an absolute standard but a relative one: it is measured against the role, the risk of the use and the context of deployment. A person who merely uses an AI system passively needs a different level of competence than a person who prepares rights- or safety-relevant decisions with AI.

From this follows a practical principle: competence cannot be determined across the board, but only on a role-by-role, risk-weighted basis. A sensible structure distinguishes along three axes:

| Axis | Example levels |
|---|---|
| AI usage | None · Low · Medium · High · Critical |
| Risk level of the use | Low · Moderate · High · Unacceptable |
| Required competence | Awareness · Foundations · Practitioner · Advanced · Responsible Use |

A role with high, risk-bearing AI usage should demonstrate a higher level of competence than a role with low usage. This is exactly the logic GARION applies: it weights readiness where the risk actually sits.

## 5.1 Examples of appropriate competence levels

The following examples are illustrative and not a binding specification; the correct classification depends on the concrete use in the organization.

| Example role | AI usage | Risk | Appropriate competence |
|---|---|---|---|
| Team member (occasional use) | Low | Low | Awareness / Foundations |
| Content creation (marketing) | High | Moderate | Practitioner (incl. source and brand safety) |
| Customer service with AI assistant | Medium | High | Practitioner |
| Recruitment / hiring | Medium | High | Practitioner to Advanced |
| Leadership with oversight role | Medium | Moderate | Responsible Use |

## 5.2 Proportionality

The regulation does not require a uniform training catalogue for everyone, but appropriate measures. For roles with low, low-risk usage, an awareness of the opportunities, limits and risks of AI may be sufficient; for roles that prepare rights- or safety-relevant decisions with AI, a considerably higher level is appropriate. This proportionate tailoring protects against two mistakes: doing too little and thereby carrying a risk — or demanding too much across the board and tying up resources without addressing the actual risk. GARION supports proportionality by assessing each role individually by usage and risk.

# 6. From intent to proof: the documentation problem

Most organizations have adopted AI organically — a tool here, a pilot there, a department quietly building it into daily work. Rarely was there a moment in which it was formally recorded who is trained, who is responsible and where the risk concentrates. The result is a gap not in capability but in proof.

This gap becomes a problem as soon as someone reviews it: a supervisory authority, a customer questionnaire, an auditor or the internal audit function. Documentation reconstructed under pressure is laborious, error-prone and unconvincing. Documentation built calmly in advance, dated and based on evidence, is cheaper to produce and considerably more robust.

> The goal is therefore not to „become AI-literate“ — many organizations already are — but to make the existing competence visible, structured and verifiable.

## 6.1 What audit-proof documentation looks like

For documentation to withstand a review, it should meet several characteristics. It is traceable (every statement can be traced back to a concrete piece of evidence), current (statuses are dated and the expiry of evidence is detected), complete to the relevant extent (the essential roles are covered), consistent (the same terms and standards across the organization) and reproducible (the same data basis leads to the same result).

These characteristics are not an end in themselves. They determine whether, in the event of a review, management can answer calmly and with evidence — or has to reconstruct under pressure. GARION is built along exactly these characteristics: deterministic calculation, end-to-end traceability to evidence, automatic status logic and an append-only audit log.

# 7. What GARION is

GARION is a local macOS and Windows application that turns scattered facts about an organizat people, roles and evidence into a single, explainable and verifiable picture of AI readiness. It is deliberately not a learning management system, not HR software and not a compliance platform: it delivers no training and claims no legal conformity, but documents that readiness exists and makes it traceable.

The model rests on four nouns:

| 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 documented, verifiable proof: a training, a certification, a policy acceptance. |

| Concept | Meaning |
|---|---|
| Readiness | The computed verdict: how well evidence covers the roles, weighted by risk. |

Four properties characterize the way it works: GARION is local (all data stays on the device, no cloud, no account), deterministic (the same inputs always yield the same number — AI never computes the score), verifiable (every figure is traceable to evidence, every change is logged) and bilingual (interface, reports and help in German and English).

# 8. Architecture and how it works

GARION is a single-process application without a server and without a cloud database. An interface presents the content; a native core handles data storage, calculations, the audit log and the generation of reports. The only optional network call is an explicitly approved request to an AI service that exclusively formulates explanatory text.

The order in which GARION works follows a clear doctrine: Readiness → Risks → Gaps → Recommendations → Details. You record roles and evidence; a deterministic engine computes the readiness score from them; a second engine derives prioritized gaps; and the reporting layer assembles everything into verifiable artifacts. All core functions work fully offline; the optional AI is additive and removable without losing a number, a gap or a report.

## 8.1 Local-first in practice

„Local-first“ is more than a technical detail. It means that the readiness data — the statements about who in the organization is how competent — does not leave the device. There is exactly one database file; a backup is a copy of that file, a restore replaces it. There is no server that could fail and no account that could be compromised. For an organization that processes sensitive personnel and governance data, this architecture is the simplest form of data minimization: what is never transmitted does not have to be secured either.

# 9. The role model

Roles describe which activities carry which AI risk. They are the basis for GARION weighting readiness where the risk sits. So that an organization does not start in front of an empty system, GARION ships four bilingual template packs with 24 predefined roles (General Business, L&D/HR, Sales & Marketing, Management). Each role is an editable template — the organization adapts it to its reality.

Each role's profile comprises AI usage, risk level, required competence, the required evidence types as well as notes on typical risks and likely gaps. Higher risk increases the role's weight in the overall score. Template roles are read-only; duplicating creates an editable copy.

### 9.1 The four template packs

The included packs cover typical organizational structures and can be freely adapted or extended:

| Pack | Included roles (examples) |
|---|---|
| General Business | Administrative Assistant, Office Manager, Customer Service, Project Coordinator, Team Member, Team Lead |
| L&D / HR | HR Manager, Recruiter, L&D Manager, Instructional Designer, Training Coordinator, HR Business Partner |
| Sales & Marketing | Marketing Manager, Content Manager, Sales Representative, Sales Manager, Account Manager, Business Development |
| Management & Leadership | Managing director / CEO and other leadership functions with organization-wide effect |

In total 24 predefined roles ensure that an organization does not start in front of an empty system, but immediately receives a realistic, adaptable starting picture.

## 10. Evidence: types, status and audit

Evidence is a documented proof of readiness — for example a completed training, not the training itself. GARION recognizes a controlled vocabulary of ten evidence types: AI literacy training, certification, assessment, policy acceptance, manager confirmation, AI tool enablement, LMS completion, workshop, document and other. This keeps records consistent and analyzable across the organization.

Each piece of evidence automatically and deterministically receives a status that is recomputed on every access:

| Condition | Status |
|---|---|
| No issue date | Pending |
| No expiry or expiry more than 30 days away | Valid |
| Expiry within 30 days | Expiring |
| Expiry in the past | Expired |

Every action — Created, Updated, Assigned, Archived, Expired — is recorded both organization-wide and in the evidence item's own history. When a piece of evidence moves to „Expired“, this is logged automatically. In this way the readiness picture stays honest without manual monitoring: it shows decay before it becomes a finding.

# 11. The readiness score

The readiness score is a single number from 0 to 100. It is deterministic, decomposable, named in four bands and accompanied by a confidence level. The entire calculation can be opened in a „How was this calculated?“ view — with the sub-scores, their weights, the risk weighting per role and the confidence breakdown.

## 11.1 The eight components and their weights

| # | Component | Weight | Measures |
|---|---|---|---|
| C1 | Evidence Coverage | 22 % | Share of required evidence present and valid. |
| C2 | Role Coverage | 18 % | Share of roles meeting their threshold. |
| C3 | Risk Exposure (inverse) | 20 % | Uncovered readiness on roles with higher risk. |
| C4 | Training Freshness | 12 % | Share of fresh (not stale) training evidence. |
| C5 | Responsibility Coverage | 10 % | Roles and key evidence with an assigned owner. |
| C6 | Governance Signals | 8 % | Recency of the assessment and structural completeness. |
| C7 | Policy Acceptance | 6 % | Share of required policy acceptances obtained. |
| C8 | Assessment Completion | 4 % | Whether a current assessment exists. |

Evidence Coverage (C1) and risk-weighted exposure (C3) together carry 42 %, because the core is „evidence- and risk-oriented“. Governance, policy and assessment (C6–C8) are real but secondary signals and deliberately weighted low.

## 11.2 Risk weighting

Within the role-related components, each role contributes proportionally to a risk weight. A high-risk role without competence evidence lowers the score about three times more than a role with low risk. This mechanism reflects the risk-based spirit of the EU AI Act without claiming legal conformity.

## 11.3 Confidence

Confidence is separate from the score and answers how much you should trust the number. It results from four factors: data completeness (40 %), evidence freshness (25 %), source quality (20 %) and recency of input (15 %), mapped to Low, Medium or High. A high score from few data points is not the same as a high score from full coverage; confidence prevents false comfort.

### 11.4 The four bands

| 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. |

The bands are advisory descriptions, not conformity verdicts. The interface deliberately never uses the word „compliant“.

### 11.5 How the calculation works at its core

The overall score is the weighted average of the eight component sub-scores (each 0–100), with the role-related components risk-weighted. By way of example: Training Freshness is the share of fresh among all training evidence (expiring counts proportionally); Policy Acceptance is the share of obtained among required acceptances, risk-weighted by role; the Governance Signals blend the recency of the last assessment with the structural completeness of the data.

Important for practice: imported evidence is scored identically to manually entered evidence once committed; source quality is judged by the evidence type, not by how it arrived. Importing existing data, however, raises data completeness — and thereby lifts both the score and the confidence. The weights and thresholds are stored as versioned constants, so an adjustment remains traceable and reproducible.

## 12. The gap engine

A number alone is not enough; what matters is what to do. The gap engine is a deterministic rules engine that runs after readiness and produces a list sorted by impact. It applies eight rules; the base severity is escalated by the role's risk weight.

| Rule | Detects |
|---|---|
| GR1 | Missing evidence: a required evidence type has no valid coverage. |
| GR2 | Expired evidence still relied upon. |
| GR3 | High-risk role below the required competence. |
| GR4 | Missing responsibility: role or key evidence without an owner. |
| GR5 | Missing policy acceptance. |
| GR6 | Stale training (past the staleness horizon). |
| GR7 | Incomplete role coverage (unmapped or below threshold). |
| GR8 | Governance blind spot (no assessment for 90+ days or risk levels unset). |

Each gap carries a severity (Low · Medium · High · Critical), an uplift estimate (how much readiness could improve once fixed — an estimate, not a promise) and a concrete, pre-written action. In this way a diagnosis immediately becomes a prioritized, trackable action plan; the lifecycle of each gap (open, in progress, resolved) is steerable.

### 12.1 Typical measures per gap type

| Gap | Typical measure |
|---|---|
| GR1 Missing evidence | Record the required evidence or run the associated training. |
| GR2 Expired evidence | Renew the evidence and enter the new issue date. |
| GR3 Role below required competence | Targeted qualification for the affected role until the target level is proven. |
| GR4 Missing responsibility | Assign a responsible person (owner) to the role or evidence. |
| GR5 Missing policy acceptance | Obtain and document acceptance of the relevant policy. |
| GR6 Stale training | Refresh the training and update the evidence. |
| GR7 Incomplete role coverage | Map roles completely and add missing profiles. |
| GR8 Governance blind spot | Run a fresh assessment and set missing risk levels. |

Because GARION sorts the gaps by a priority score (not by severity alone), the most effective measure stands first — the organization works from top to bottom and sees the progress after each recalculation.

## 13. The EU AI Act Readiness Check

GARION includes its own, purely deterministic readiness check. It checks in a structured way whether a report carries sufficient readiness evidence and documentation quality. It is explicitly not a conformity certification; the naming is binding „Readiness Check“, never „Compliance Certification“. Each area is scored deterministically (covered, partial, missing) and aggregated into an overall status: Strong · Partial · Weak · Not enough data.

Ten areas are checked:

- AI literacy documented
- Role-based AI usage documented
- Required evidence assigned
- Policy acceptance documented
- Responsibility/owner assigned
- Risk level documented

- Readiness gaps identified
- Recommended actions included
- Audit trail available
- Non-legal notice included

The check moves the question from the impermissible („Are we compliant?“) to the answerable („Is our documentation strong enough to defend?“). The legal assessment remains expressly reserved for a competent body.

# 14. Reports and exports

The reporting layer produces verifiable artifacts, each with a provenance and disclaimer block. Every report is available in German or English and works fully with AI disabled (deterministic text building blocks instead of an AI narrative).

| Artifact | Content |
|---|---|
| Executive One-Pager (PDF) | Single-page management overview: score, confidence, top-3 risks/gaps/actions, EU AI Act Readiness status. |
| AI-Readiness Report | The complete, explainable scorecard. |
| Audit-Readiness Report | Check whether a defensible report can be exported, with blocking items and fixes. |
| Role reports | Detailed per function. |
| XLSX matrices / CSV | Role, evidence, gap, department and audit tables for detail work and further processing. |
| .garion backup | Portable package for backup and handover; contains no secrets (no API key). |

Every report carries the notice that GARION does not certify legal conformity and does not replace legal, regulatory or professional advice.

## 14.1 Which report for what

The reports serve different purposes. The Executive One-Pager gives leadership a condensed status overview for reporting and decisions. The AI-Readiness Report delivers the complete, explainable basis when a technical or compliance function wants to follow the derivation. The Audit-Readiness Report answers whether the documentation is currently verifiable and names the blocking items. The matrices and CSV exports support operational detail work, and the .garion backup serves for safekeeping and a clean handover.

It makes sense to use the Audit-Readiness Report as an internal early-warning system: it shows what would currently stand in the way of a review, and thereby makes preparation plannable — long before an external enquiry arrives.

## 15. Data protection and security

In a regulatory environment, the confidentiality of the data is central. GARION is built so that the readiness data does not leave the device:

- **Local-first.** The application runs entirely on the device; the core functions work without a network.
- **Keychain.** An optional AI key lives only in the operating system keychain — never in the database, configuration or logs.
- **No content to the AI.** No file contents or evidence are sent to the AI; aggregation and redaction apply by default.
- **Deterministic core.** Scores and statuses are computed by fixed rules; the AI is additive and removable.
- **Append-only audit log.** Relevant actions and every AI call are captured and exportable.

The exact location of the local database is shown in the settings and under „About GARION“. A backup is a copy of that file; a restore replaces it.

## 16. Recommended workflow

The following workflow takes an organization from an empty state to a robust, documented readiness picture. A first, meaningful assessment is achievable in under 20 minutes.

1. **Create the organization** and choose the working language.
2. **Install the template packs** and trim to the roles that actually exist.
3. **Set AI usage and risk per role** — classify the rights- or safety-relevant roles realistically in particular.
4. **Record or import evidence** (CSV/XLSX) — e.g. completed trainings, certifications, policy acceptances.
5. **Compute readiness** and save a first snapshot.
6. **Work the gaps in priority order** — first those with the highest uplift; record each remediation.
7. **Export reports** — Executive One-Pager for leadership, Audit-Readiness Report for internal audit.
8. **Set a re-assessment cadence** (e.g. quarterly) to build a trend.

## 17. Roles and responsibilities in the company

AI readiness is a cross-cutting task. A clear allocation prevents it from getting lost between departments. The following split has proven itself in practice and can be adapted to the size of the organization:

| Function | Contribution |
|---|---|
| Executive management / leadership | Bears overall responsibility; uses the score and One-Pager for reporting and decisions. |
| AI / compliance officer | Maintains the readiness picture, defines roles and risk, steers gap remediation. |
| Department / team leads | Confirm the real AI usage of the roles and name owners per area. |
| HR / L&D | Supplies training and competence evidence; links trainings to roles. |
| IT / data protection | Responsible for device, backup and the local data store; checks the data-protection configuration. |

It is important that every role and every key piece of evidence has an owner. GARION makes missing responsibilities visible via rule GR4.

## 18. Governance over time

Competence is not a one-time event: certifications expire, policies are superseded, teams and AI usage change. A one-time proof therefore inevitably ages. GARION addresses this in two ways.

First, the automatic expiry management: evidence moves to the status „Expiring“ 30 days before expiry and thus gives an early prompt to renew. Second, the snapshot mechanism: every recalculation saves a state, and every report shows the change versus the last measurement. In this way a trend emerges that demonstrates progress over time — exactly what an ongoing, documented „best efforts“ approach looks like in practice.

In practice, a fixed cadence with clear responsibility is recommended: one responsible person, a recurring appointment (e.g. quarterly or upon material changes in AI usage) and a short, always identical routine — record new evidence, renew expired evidence, adapt roles where usage changes, recalculate, export the report. Triggers for an unscheduled re-assessment are the introduction of a new AI system, organizational changes and regulatory updates. Because the effort per cycle is low, the documentation stays current rather than ageing between two large projects.

## 19. Limits: what GARION does not do

An honest delimitation is part of a robust tool. GARION expressly does not do the following:

- **No legal advice and no conformity certification.** GARION assesses the quality of evidence and documentation, not legal conformity.
- **No substantive assessment of AI systems.** GARION checks organizational readiness, not the technical safety or functioning of individual systems.
- **No substitute for training.** GARION documents competence evidence but does not itself deliver content.
- **No automatic decisions.** The AI only formulates explanations; the final interpretation is confirmed by a human.

These limits are not a deficiency but the precondition for every statement of the tool remaining defensible: GARION never promises more than it can keep.

### 19.1 Summary

The EU AI Act primarily requires companies not so much to become AI-literate, but to be able to prove that the right roles hold the right, risk-appropriate competence — and that this is documented. This obligation has been active since February 2025 and enforceable from August 2026. Whoever builds the evidence calmly, dated and traceably in advance is, in the event of a review, proven rather than under pressure. GARION provides the structured, local and verifiable factual basis for this: an explainable readiness picture, prioritized gaps with measures and verifiable reports — without making legal statements and without data leaving the device. The legal assessment remains reserved for a competent body; the factual basis for it is created by the organization itself.

## 20. Implementation checklist

A compact list that takes an organization from introduction to a maintained readiness program:

- Organization created, working language chosen.
- Roles recorded completely, AI usage and risk set realistically.
- Required evidence types defined per role.
- Existing evidence recorded or imported; sources clean.
- Owners assigned for roles and key evidence.
- Policies stored and their acceptance documented.
- First readiness snapshot saved (baseline).
- Gaps prioritized by uplift and in progress.
- Executive One-Pager and Audit-Readiness Report exported.
- Re-assessment cadence agreed (e.g. quarterly).

- Data backup set up as a .garion package.

## 21. Frequently asked questions

| Question | Answer |
|---|---|
| Do we need an account or internet? | No. GARION is local-first; core operation needs neither an account nor an internet connection. Only the optional AI explanation needs a provider key and a connection. |
| Does the AI change our values? | No. All core values are deterministic. The AI only formulates explanations in plain language and can be switched off. |
| Does GARION certify conformity? | No. GARION delivers a deterministic readiness view and does not replace a legal or professional review. |
| Where is our data? | In a local database on the device. The exact path is shown in the settings and under „About GARION“. |
| Is content sent to the AI? | No. No file contents or evidence are transmitted; aggregation and redaction apply by default. |
| How often should we re-assess? | A quarterly re-assessment has proven useful; it keeps the documentation current and builds a trend. |
| Which platform is supported? | macOS on Apple Silicon and Windows 10/11 (64-bit). No internet is required for core opera |

Note: GARION is available in two editions with an identical core. The Business Edition assesses a single organization; the Consultant Edition manages several strictly separated clients (its own database per client, no cross-client access). For a single company, the Business Edition is the right choice.

## 22. Glossary

| Term | Meaning |
|---|---|
| Readiness score | A value from 0–100 for organizational AI readiness, deterministic and risk-weighted. |
| Confidence | A measure of how much to trust the score (data completeness, recency, source quality). |
| Evidence | A documented, verifiable proof of readiness. |
| Coverage | Share of required evidence that is present and valid. |
| Gap | A deterministically detected shortfall with severity, uplift and a measure. |
| Uplift | Estimated readiness improvement after fixing a gap — an estimate. |
| Band | Named range of a score: At Risk / Building / Developing / Ready. |
| Snapshot | A saved readiness and confidence value at a point in time; enables a trend. |
| Audit log | Append-only record of relevant changes and every AI call. |
| Deployer | Whoever uses an AI system under its own responsibility. |

# 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 (in German: EU-KI-Verordnung) 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 organization

Getting started becomes easier and the view becomes wider: existing organizational data enters GARION through a controlled import instead of manual entry — pseudonymous, validated and auditable. And alongside AI literacy documentation under the EU AI Act, GARION now also answers leadership’s cyber governance question: which risks are open, which treatments are running, which evidence exists, which deadlines are due — at a glance, locally, without cloud.

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