PLATFORM DOCUMENTATION

# The E-ARI Handbook

A complete guide to the platform — what it does, how it works, and why it's built the way it is.

VERSION 1.0 · MAY 2026

E-ARI

## Foreword

Most “AI readiness” tools are spreadsheets in disguise. They ask a few questions, draw a radar chart, and hand you a PDF that ages out the moment your CFO reads it.

E-ARI was built to do something harder: turn the messy, regulated, fast-moving reality of enterprise AI adoption into a single living workspace where measurement, evidence, and compliance compound over time. One number on a dashboard is easy. A defensible audit trail that survives a Notified Body review three years from now is not.

This document explains every part of the platform — the assessment engine that grades your readiness, the six AI agents that turn scores into narrative, the compliance autopilot that maps your AI systems to the EU AI Act, and the monitoring layer that watches for new gaps while you sleep.

Read it cover-to-cover or jump to the section you need. There’s a glossary at the end if a term feels unfamiliar.

---

## Table of contents


1. What E-ARI is
2. The problem we’re solving
3. How E-ARI is structured
4. Layer 1 — The Assessment Engine
5. Layer 2 — The Six AI Agents
6. Layer 3 — The Compliance Autopilot
7. Layer 4 — Continuous Monitoring
8. The four-step user journey
9. Methodology & scoring
10. Privacy, security, and data handling
11. Plans & pricing
12. Who this is for
13. Roadmap
14. Glossary


---

# 1. What E-ARI is

E-ARI — short for **Enterprise AI Readiness Index** — is an EU-AI-Act-aware platform that helps an organisation answer four questions, in order:

| QUESTION | E-ARI ANSWER |
|---|---|
| **How mature are we at AI adoption?** | An 8-pillar readiness score, benchmarked against your sector. |
| **Which obligations apply to our AI systems?** | An automatic risk classification per AI system (Article 5 / 6 / 7 + Annex III), with a per-system obligation list. |
| **Do we have the evidence to back it up?** | A typed evidence vault with automatic clause extraction, mapping each piece of evidence to the obligations it supports. |
| **What changed this week?** | Daily regulatory scan + monitoring plan + automated reminders before attestation deadlines. |

The output is not a one-shot PDF. It is a living workspace: scores, FRIAs, Technical Files, gap reports, and obligation coverage all stay versioned, exportable, and audit-ready.

---

# 2. The problem we're solving

## 2.1 The compliance landscape changed in 2024

The **EU AI Act** entered into force on 1 August 2024 with phased application:

- **2 February 2025** — Articles 5 (prohibited practices) and 4 (AI literacy) apply.
- **2 August 2025** — Governance, GPAI obligations, penalties, and notified body designations.
- **2 August 2026** — High-risk system obligations apply in full (Annex III + Annex I systems).
- **2 August 2027** — Embedded high-risk systems in regulated products (Annex I).

Penalties scale to **€35M or 7% of global turnover** for prohibited-practice violations. There is no "we'll get to it next quarter" version of this.

### 2.2 Existing tools fail in three predictable ways

| FAILURE MODE | WHAT IT LOOKS LIKE | WHAT IT COSTS YOU |
|---|---|---|
| **Spreadsheet maturity models** | Four colours on a tab, copied between teams every quarter. | No version control, no defensibility, no link to actual evidence. |
| **Big-Four consultancy projects** | A 200-page report that's accurate the day it ships. | High cost, low cadence — by the time the next assessment runs, the regulatory landscape has moved. |
| **Generic GRC platforms** | Compliance modules with a generic "AI" checkbox bolted on. | No native handling of FRIA (Article 27) or Annex IV Technical Files. No risk classifier. No clause-level evidence mapping. |

### 2.3 What E-ARI optimises for

Three properties that we treat as non-negotiable:

1. **Defensibility** — every recommendation, every classification, every gap is traceable to a clause in a regulation, a sentence in a piece of evidence, or a question in the assessment. No black-box LLM verdicts.
2. **Compounding** — work done in one place feeds the next. Your assessment baseline classifies your AI systems. Evidence uploaded for one obligation auto-maps to others. Closing a gap on one system updates monitoring across the portfolio.
3. **Surface area discipline** — we deliberately do *not* try to be a project management tool, a GRC giant, or a model registry. E-ARI sits between your strategy team and your regulators.

---

## 3. How E-ARI is structured

The platform is built as four cooperating layers. Each layer has its own data model, its own UI surfaces, and its own integration points — but they share a single source of truth.


A baseline **Assessment** is the seed crystal. From there, every AI system you register gets bound to that baseline so its compliance profile inherits the organisational context (sector, size, governance maturity) automatically.

---

# 4. Layer 1 — The Assessment Engine

## 4.1 The 8 pillars

We score readiness across **eight pillars**, with weights derived from research literature on AI adoption success factors. The weights sum to 1.00.

| # | PILLAR | WEIGHT | WHAT IT MEASURES |
|---|---|---|---|
| 1 | **Strategy & Vision** | 15% | Formal AI strategy, executive sponsorship, ROI measurement, multi-year roadmap. |
| 2 | **Data & Infrastructure** | 15% | Data quality, accessibility, governance, ETL/MLOps maturity, compute. |
| 3 | **Technology & Tools** | 12% | ML platforms, model lifecycle tooling, vendor diversity, cloud architecture. |
| 4 | **Talent & Skills** | 13% | Hiring pipeline, internal training, role definitions, succession planning. |
| 5 | **Governance & Ethics** | 15% | Ethics review boards, policy completeness, bias auditing, accountability. |
| 6 | **Culture & Change** | 10% | Adoption psychology, change management, communication, resistance handling. |
| 7 | **Process & Operations** | 10% | Workflow integration, process re-engineering, feedback loops, SLA management. |
| 8 | **Security & Compliance** | 10% | Model security, privacy controls, regulatory alignment, audit readiness. |

Each pillar contains exactly **5 questions** on a Likert scale (1–5). 40 questions total. Median completion time: ~15 minutes.

## 4.2 Maturity bands

Once weighted, the overall score (0–100) maps to one of four bands:

| BAND | RANGE | INTERPRETATION |
|---|---|---|
| 🔴 **Laggard** | 0–25 | Minimal foundational elements. AI initiatives are ad-hoc or absent. |
| 🟠 **Follower** | 26–50 | Early-stage readiness. Some initiatives but lacking cohesion or alignment. |
| 🔵 **Chaser** | 51–75 | Progressing readiness. Foundations in place, active investment underway. |
| 🟢 **Pacesetter** | 76–100 | Advanced readiness. Well-positioned for AI-driven competitive advantage. |

Bands are deliberately broader than a percentile — they're meant to drive *behaviour* (where to invest), not vanity metrics.

### 4.3 What you get back

A completed assessment produces a structured `Assessment` record:

- **Overall score** (0–100, weighted)
- **Per-pillar scores** with question-level breakdowns
- **Maturity band** with a plain-English interpretation
- **Sector benchmark comparison** (your scores vs. curated industry averages)
- **Strength / weakness matrix** — top three pillars and bottom three
- **Initial regulatory mapping** — pillars that align with applicable regulations (GDPR, EU AI Act, ISO 42001, NIST AI RMF)

This record is the seed for everything else.

---

## 5. Layer 2 — The Six AI Agents

When an assessment finishes, an orchestrator runs a six-agent pipeline. Each agent has a single job, and the outputs of earlier agents feed the inputs of later ones.

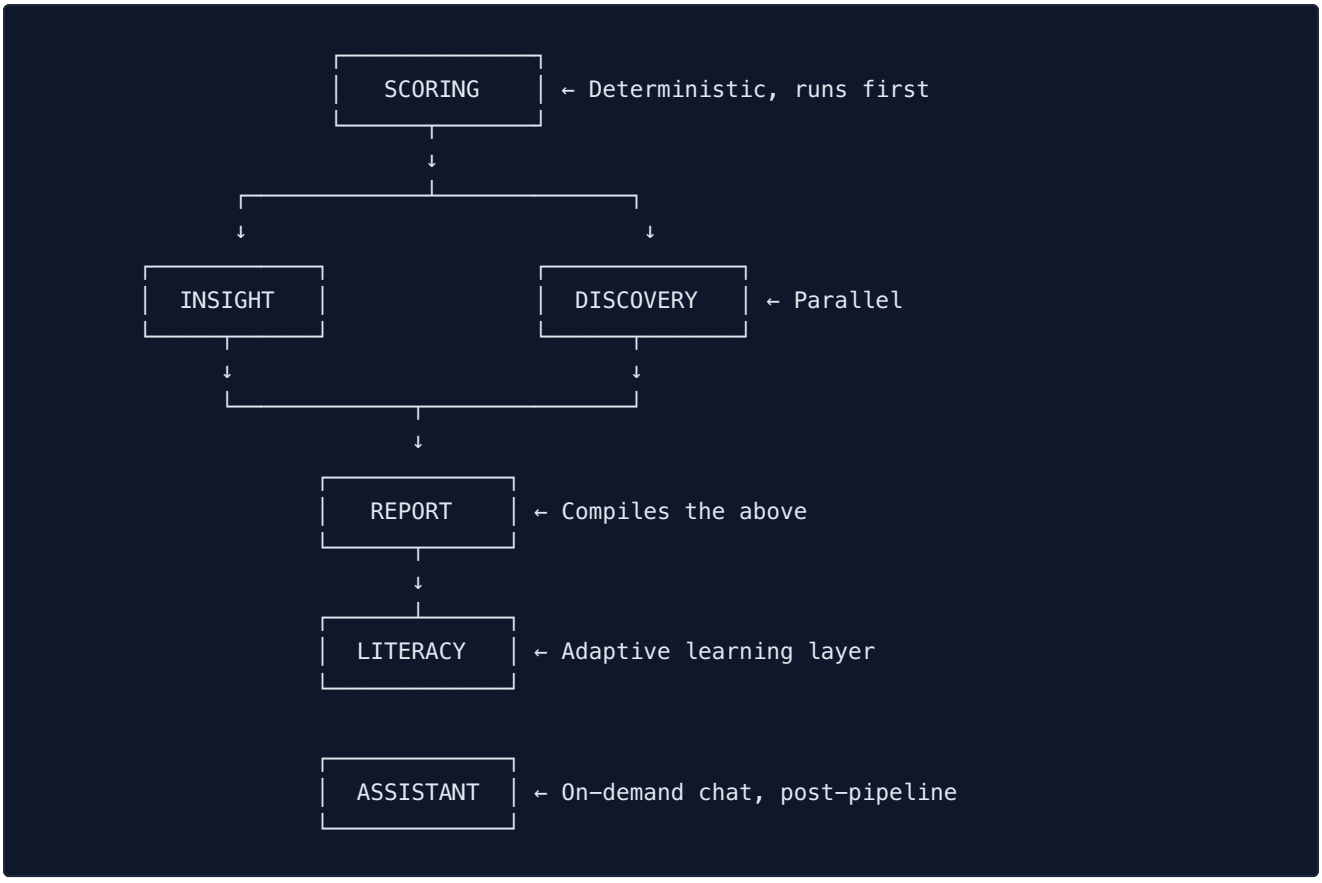

## 5.1 Scoring Agent

**Deterministic**, no LLM involvement, runs first. Computes:

- Per-pillar weighted average from the Likert responses
- Overall score = Σ (pillar_score × pillar_weight)
- Sector-adjusted percentile (using the curated benchmark library)
- Confidence interval (function of question completion rate and response variance)

The scoring math is open and reproducible. It is the only deterministic step in the pipeline — every other agent grounds its outputs against this scoring result, which means LLM hallucinations cannot move your score.

## 5.2 Insight Agent

LLM-powered, generates **narrative insights grounded in the scoring result**. For each pillar:

- A 2–3 sentence interpretation of the score in business language
- Two specific risks if the score stays where it is
- Two specific opportunities unlocked by reaching the next maturity band

Outputs are run through a grounding verifier — claims that don't tie back to a numeric score in the scoring result are stripped out before display.

## 5.3 Discovery Agent

Runs **in parallel with Insight** because their inputs don't overlap. Discovery answers the question *"what does your organisation actually look like to the outside world?"*:

- Auto-classifies your sector via a 30-token LLM call (no cross-industry pollution)
- Pulls public signals via Tavily Search and Extract (recent news, regulatory filings, public hiring patterns)
- Synthesises a contextual brief: sector trends, peer announcements, observed regulatory pressure
- Citations are inline `[S1]`, `[S2]` — no source, no claim

If the public signal is weak or untrustworthy, Discovery declines to synthesise rather than fabricating a "medium-confidence" reading.

## 5.4 Report Agent

Compiles everything above into a structured executive document:

- Executive summary (one page, weighted by your sector context)
- Pillar-by-pillar deep dive with insights and risks

- Top three recommended actions, each tied to specific pillars and obligations
- Benchmarking section with the sector context from Discovery
- Appendix: full question responses, methodology notes, regulatory mapping

Available as an in-app view, a `.docx` download, and a print-optimised PDF.

### 5.5 Literacy Agent

Adaptive, runs after Report. Looks at *your* weakest pillars and assembles a personalised mini-curriculum:

- Targeted articles from the Literacy Hub
- Quizzes calibrated to your maturity band
- Recommended learning paths (e.g. "From Follower to Chaser on Governance" — 4 modules, ~3 hours)

Unlike a static knowledge base, Literacy never recommends content for pillars you're already strong in. It earns its position in the pipeline by being personalised.

### 5.6 Assistant Agent

On-demand, post-pipeline. A grounded chat interface that has read-only access to your assessment, your AI systems, your evidence, and your obligation coverage.

- Asks: *"What's blocking us from being a Chaser?"*
- Refuses to answer questions outside its data scope (no general AI advice — only what's grounded in your platform state)
- All citations link back to specific records in your account

---

## 6. Layer 3 — The Compliance Autopilot

This is where E-ARI stops being a measurement tool and starts being a defensibility engine. The Compliance Autopilot maps your AI systems to the **EU AI Act** and produces the artefacts a regulator (or your own internal audit team) will ask for.

### 6.1 The AI System registry

Every AI system you operate or deploy gets registered with:

- **Name & purpose** — plain-English description
- **Deployer role** — provider, deployer, importer, distributor (matters for which obligations apply)
- **Sector & populations affected** — drives risk classification logic

- **Link to your readiness baseline** — inherits your organisational context automatically

A single account can register unlimited systems on the Growth and Enterprise tiers.

## 6.2 The risk classifier

For each registered system, a classifier runs against the Act's risk taxonomy:

| RISK TIER | TRIGGER | WHAT WE GENERATE |
|---|---|---|
| 🚫 **Prohibited** (Art. 5) | Subliminal/exploitative practices, social scoring, real-time biometric ID in public spaces | Cease-and-desist guidance + cited articles |
| 🔴 **High-risk** (Art. 6 + Annex III) | Safety components, biometrics, education, employment, essential services, law enforcement, etc. | Full Annex IV Technical File + FRIA + post-market monitoring plan |
| 🟡 **Limited risk** (Art. 50 / 52 / 53) | Chatbots, deepfakes, emotion recognition, biometric categorisation | Transparency obligation kit (notice templates, disclosure copy) |
| 🟢 **Minimal risk** | Everything else | Lightweight obligation list (literacy, voluntary code adherence) |

Each classification ships with a **rationale**: the specific articles that triggered the tier, the public guidance interpreting them, and the residual ambiguity (so your DPO can override if the regulator's view evolves).

## 6.3 The evidence vault

A typed object store for compliance evidence:

- Contracts, DPIAs, model cards, dataset documentation, training records, security audits, fairness reports, etc.
- Org-level evidence (applies to all systems) and per-system evidence
- Versioned — replacing a piece of evidence creates a new version, the old one stays linked to historical FRIAs
- Stored in **Vercel Blob** with private ACL; presigned URLs for download (~5 min TTL)

When you upload, two things happen automatically:

1. **Classification** — an LLM categorises the document type (DPIA, model card, audit report, etc.)
2. **Clause extraction** — the document is chunked, and an LLM identifies which clauses or sentences map to which Article of the AI Act. These mappings are stored as `EvidenceClause` records and become reusable across obligations.

A DPIA you upload for one system can satisfy parts of three different obligations across two systems automatically.

## 6.4 The FRIA generator (Article 27)

For high-risk systems, the **Fundamental Rights Impact Assessment** is mandatory.

E-ARI's generator:

1. Pulls your AI system metadata (purpose, populations affected, deployer role)
2. Pulls the relevant evidence clauses already mapped to fundamental rights articles
3. Generates a structured FRIA covering: process description, time periods, frequency, categories of natural persons affected, specific risks of harm, governance arrangements, complaint mechanisms
4. Surfaces gaps inline (e.g. *"No evidence found for human oversight design — Article 14 requirement"*)
5. Exports as PDF and `.docx` with embedded source citations

You can finalise a FRIA only when all required sections have backing evidence. Half-finished FRIAs are saved as drafts.

## 6.5 The Technical File generator (Annex IV)

The Annex IV Technical File is the formal document that a Notified Body or supervisory authority will request first.

E-ARI generates the full structure required by Annex IV:

1. General description of the AI system
2. Detailed description (development methods, third-party tools, training datasets)
3. Information on monitoring, functioning, and control
4. Risk management system documentation
5. Lifecycle changes log
6. Performance metrics, accuracy thresholds, robustness tests
7. List of harmonised standards applied
8. Declaration of conformity (template)
9. Post-market monitoring plan

Each section pulls from your evidence vault and your monitoring plan. Sections without backing evidence are marked clearly so you know exactly what's missing.

### 6.6 The Gap Radar

A live view of what's *missing*:

- Per-obligation: total clauses required vs. clauses currently evidenced
- Severity-weighted (a gap on Article 9 risk management ≫ a gap on Article 53 GPAI logging for a non-GPAI system)
- Aging — gaps that have been open for >30 days surface to the top
- Each gap has a one-click "Resolve" workflow that prompts for the specific evidence type needed

### 6.7 The Submission Pack

A single `.zip` export, ready to hand to a Notified Body, an auditor, or your own legal team:

- Annex IV Technical File (PDF)
- FRIA (PDF, if high-risk)
- All cited evidence files
- Obligation coverage report
- Monitoring plan
- Manifest with SHA-256 checksums for tamper evidence

---

## 7. Layer 4 — Continuous Monitoring

Compliance is not a project. It's a process that runs forever.

### 7.1 Pulse checks

Lightweight monthly check-ins targeting the questions most sensitive to drift (governance, security, talent). Takes 3–5 minutes. Generates a delta report against your last full assessment.

### 7.2 Daily regulatory scanner

A daily cron job ( `/api/cron/compliance-monitoring` ) that:

1. Scans the EU AI Act regulatory feed for new guidance, delegated acts, harmonised standards, supervisory authority opinions
2. Cross-references each item against your registered AI systems
3. Files relevant items into your **Compliance Inbox**
4. Triggers a re-classification if the change is material

### 7.3 Attestation calendar

Annex IV requires periodic attestations for high-risk systems. The platform tracks:

- Annual conformity reassessment due dates
- DPIA refresh windows (every 24 months for evolving systems)
- Post-market monitoring report deadlines
- Internal review cycles (configurable per system)

You get email reminders 60, 30, 7, and 1 day before each deadline.

### 7.4 Email & in-app notifications

All notifications are delivered via the unified template system:

- Daily compliance digest (opt-in, only when there's something new)
- Real-time alerts on critical regulatory changes
- Weekly executive summary (for users tagged as executives)

Unsubscribe and preference management at `/portal/preferences`.

---

## 8. The four-step user journey

Every user experiences the platform as four sequential steps. The dashboard makes this explicit with a progression banner.

### Step 1 — Assess

Run the 8-pillar assessment. ~15 minutes. Produces a baseline that anchors all downstream compliance work.

### Step 2 — Verify

Register your AI systems, upload evidence, let the classifier run. The platform tells you exactly which obligations apply to each system.

### Step 3 — Comply

Generate FRIAs and Technical Files. Close gaps. Export Submission Packs when audited.

### Step 4 — Monitor

Subscribe to the regulatory scanner. Honour the attestation calendar. Run quarterly pulse checks to catch drift early.

The progression is visualised in your portal as a single horizontal stepper — `Assess → Verify → Comply → Monitor` — with the current step highlighted and a contextual CTA pointing to the next concrete action.

---

# 9. Methodology & scoring

## 9.1 Scoring math

For each pillar $p$ with questions $q_1 \ldots q_5$ answered on a 1–5 Likert scale:

```
pillar_score(p) = (Σ qᵢ − 5) / 20 × 100
```

This normalises the raw 5–25 sum to a 0–100 scale.

The overall score is the weighted average:

```
overall_score = Σ (pillar_score(p) × weight(p))
```

with weights summing to 1.00. The result is rounded to the nearest integer for display but stored at full precision.

## 9.2 Confidence interval

A scoring result includes a confidence band derived from:

- **Question completion rate** — partial assessments get wider bands
- **Response variance** — extreme highs and lows in adjacent pillars suggest survey fatigue
- **Free-text consistency** — when an elaboration contradicts the Likert score, confidence drops

The band is shown in the UI as a range (e.g. *"68 ± 4"*) and is preserved in exports.

## 9.3 Sector benchmarks

Benchmarks are curated from public research (industry association reports, academic studies, regulator-published statistics). They are explicitly labelled in the UI:

> ⚠️ *Sector benchmarks are AI-estimated and intended for directional guidance only. Actual sector performance may vary.*

We never claim a benchmark is current — we cite the year of the underlying data.

### 9.4 Why we don't use a single "trust score"

Some platforms compress everything into a single 0–100 "trustworthiness" number. We don't. A single number obscures the only useful signal: *which dimension is failing*. An organisation can be a Pacesetter on Strategy and a Laggard on Security simultaneously — that's actionable. Averaging it to "Chaser, 62" is not.

---

## 10. Privacy, security, and data handling

### 10.1 What we store

| CATEGORY | EXAMPLES | WHERE |
|---|---|---|
| **Account info** | Name, email, organisation, role, password hash | PostgreSQL (Supabase) |
| **Assessment responses** | Likert answers, free-text elaborations | PostgreSQL |
| **AI system metadata** | System name, purpose, sector, classification | PostgreSQL |
| **Evidence documents** | Uploaded files (DPIAs, contracts, model cards) | Vercel Blob (private) |
| **Generated artefacts** | FRIAs, Technical Files, reports | PostgreSQL + Vercel Blob |
| **Logs** | Compliance audit log, user actions | PostgreSQL with 90-day retention |
| **Payment** | Customer ID, subscription tier | Stripe (we never store card numbers) |

### 10.2 Encryption

- **In transit** — TLS 1.3 everywhere, HSTS preload-listed
- **At rest** — AES-256 (provided by Supabase and Vercel Blob)
- **Passwords** — bcrypt with cost factor 12

### 10.3 Authentication

- Email + password (NextAuth credentials provider)
- Google OAuth (NextAuth Google provider)
- SSO/SAML on Enterprise tier (Okta, Azure AD, Google Workspace)

Sessions use JWT with HttpOnly, Secure, SameSite=Lax cookies.

### 10.4 Data residency

EU residency on Enterprise tier (Supabase EU region + Vercel Blob EU region). Default tier uses US-East with cross-border data transfer addendum (SCCs in the DPA).

### 10.5 Data deletion

- Account deletion is self-service from `/portal/preferences`
- Deletion is hard (not soft) — within 30 days, all account data including evidence files, assessments, FRIAs, and logs are purged
- Audit log retention is independent and purged after 90 days regardless of account state

### 10.6 GDPR posture

- DPA available on request (`hello@e-ari.com`)
- Sub-processor list maintained at `/privacy#sub-processors`
- Right to erasure honoured within 30 days
- Right to portability — full account export as JSON from `/portal/preferences`

### 10.7 Third-party LLM use

We use multiple LLM providers (Gemini, GLM, NVIDIA-hosted models). All providers are configured with **zero data retention** for our API key. We do not use customer data to train any model. Free-text inputs are redacted of common PII patterns before being sent to any LLM endpoint.

---

## 11. Plans & pricing

| PLAN | PRICE | BEST FOR |
|---|---|---|
| **Starter** | Free | Solo practitioners exploring the platform |
| **Professional** | €49/month or €490/year (save 17%) | Practitioners running regular assessments |
| **Growth** | €149/month or €1,490/year (save 17%) | Scaling organisations with multiple AI systems |
| **Enterprise** | Custom | Regulated industries, multi-org deployments |

### 11.1 What's included by tier

| CAPABILITY | STARTER | PROFESSIONAL | GROWTH | ENTERPRISE |
|---|---|---|---|---|
| **Full assessments / month** | 1 | 5 | 20 | Unlimited |
| **Pulse checks / month** | 3 | 15 | 50 | Unlimited |
| **Team members** | 1 | 5 | 25 | Unlimited |
| `.docx` **reports** | 1/mo (€29 add-on) | 3 included | Unlimited | Unlimited + custom branding |
| **Literacy Hub** | Basic | Full library | Full + Learning Paths | Full + custom content |
| **Sector benchmarks** | 1 sector | 5 sectors | All sectors | All + custom benchmarks |
| **Compliance Autopilot** | — | — | ✓ | ✓ |
| **FRIA + Technical File** | — | — | ✓ | ✓ |
| **Admin portal** | — | Basic | Full | Full + SSO/SAML |
| **API access** | — | — | Read-only | Full CRUD |
| **Support** | Community | Email (48h) | Chat + quarterly review | Dedicated CSM + SLA |

### 11.2 Cancellation

Cancel any time. Access remains until end of billing period. Account auto-reverts to Starter — your data is preserved, you simply lose access to paid-tier features.

---

## 12. Who this is for

### 12.1 Three primary personas

**The AI Programme Lead** at a mid-market enterprise (~500–5,000 employees). Reports to the COO or CIO. Owns the AI strategy roadmap. Needs E-ARI to:

- Demonstrate progress to executives quarter over quarter

- Translate readiness gaps into specific investment proposals
- Show the board the company isn't behind peers

**The DPO / Head of Compliance** at any organisation deploying AI in EU jurisdictions. Needs E-ARI to:

- Auto-classify AI systems against the Act
- Maintain a living evidence vault that survives an audit
- Generate FRIAs and Technical Files without retaining a Magic Circle law firm

**The Internal Auditor** preparing for an external audit, ISO 42001 certification, or a Notified Body assessment. Needs E-ARI to:

- Produce a single Submission Pack with checksums
- Show coverage trends over time (gap closure rate)
- Provide tamper-evident logs for any artefact change

### 12.2 Who this isn't for

We're explicit about this. E-ARI is **not** the right tool if you need:

- A model registry (use MLflow, Weights & Biases, or an MLOps platform)
- A general-purpose GRC tool spanning SOC2, ISO 27001, HIPAA (use Vanta, Drata, Secureframe)
- Legal advice (E-ARI surfaces obligations; it doesn't replace counsel)
- Project management of AI delivery (use Linear, Jira, Asana)

The tighter the scope, the better the tool.

---

## 13. Roadmap

Current focus areas (subject to change):

| QUARTER | THEME | HIGHLIGHTS |
|---|---|---|
| **Q2 2026** | Connector ecosystem | Slack, Microsoft Teams, Confluence evidence import, Google Drive sync |
| **Q3 2026** | ISO 42001 alignment | Cross-mapping AI Act obligations to ISO 42001 controls |
| **Q4 2026** | Multi-jurisdiction | UK AI regulatory framework, US Executive Order, NYC AEDT compliance |
| **Q1 2027** | Notified Body workflows | Pre-conformity assessment templates, automated NB liaison packets |

Roadmap is published on the platform and updated monthly.

---

## 14. Glossary

**AI Act** — Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence. Entered into force 1 August 2024.

**Annex III** — The list of high-risk AI use cases enumerated in Annex III of the AI Act (biometrics, education, employment, essential services, law enforcement, migration, justice).

**Annex IV** — The structure of the Technical File required for high-risk AI systems.

**Article 5** — Prohibits specific AI practices: subliminal manipulation, exploitation of vulnerabilities, social scoring, real-time biometric identification in public spaces.

**Article 27** — Mandates a Fundamental Rights Impact Assessment for high-risk AI systems used by public authorities or operators of essential services.

**Article 50–53** — Transparency obligations for limited-risk systems (chatbots, deepfakes, biometric categorisation, emotion recognition).

**Conformity Assessment** — The procedure by which a high-risk AI system demonstrates compliance with AI Act requirements before being placed on the EU market.

**Deployer** — Any natural or legal person, public authority, or agency using an AI system under its authority. Distinct from Provider (the entity that develops or has the system developed).

**DPIA** — Data Protection Impact Assessment under GDPR Article 35. Often a prerequisite for FRIA.

**FRIA** — Fundamental Rights Impact Assessment under AI Act Article 27.

**GPAI** — General Purpose AI Model (foundation models). Subject to a separate obligation regime under Articles 51–55.

**Notified Body** — An organisation designated by an EU Member State to assess conformity of certain high-risk AI systems before they are placed on the market.

**Pillar** — One of E-ARI's eight readiness dimensions (Strategy, Data, Technology, Talent, Governance, Culture, Process, Security).

**Pulse Check** — A short monthly assessment targeting drift-sensitive questions.

**Submission Pack** — The single `.zip` export bundling Annex IV Technical File, FRIA, evidence, and manifest — ready for a Notified Body or auditor.

**Technical File** — The dossier required by Annex IV for high-risk AI systems. Must be kept up to date for the lifetime of the system.

---

## Appendix A — Key URLs

| SURFACE | URL |
|---|---|
| Marketing site | `https://www.e-ari.com` |
| Pricing | `https://www.e-ari.com/pricing` |
| Sign in | `https://www.e-ari.com/auth/login` |
| Portal | `https://www.e-ari.com/portal` |
| Use cases (compliance) | `https://www.e-ari.com/portal/use-cases` |
| Evidence vault | `https://www.e-ari.com/portal/evidence` |
| Privacy policy | `https://www.e-ari.com/privacy` |
| Terms of service | `https://www.e-ari.com/terms` |
| Status & health | `https://www.e-ari.com/api/health` |

# Appendix B — Contact

| TOPIC | EMAIL |
| --- | --- |
| General questions | `hello@e-ari.com` |
| Sales / Enterprise | `sales@e-ari.com` |
| Privacy / DPO | `privacy@e-ari.com` |
| Security disclosures | `security@e-ari.com` |
| Support | `support@e-ari.com` |

*E-ARI is operated by the E-ARI team. The platform is hosted on Vercel infrastructure in the United States and the European Union. This document is authoritative as of the version date above; the most current version is always served at `https://www.e-ari.com/docs/e-ari-handbook.md`.*