unorma

WHITE PAPER 1 OF 3

EU AI Act Compliance for High-Risk AI Providers

A Practitioner Framework for Technical Documentation, Conformity Assessment, and Post-Market Monitoring Under Chapter III

March 2026 · unorma.com/eu-ai-act-compliance

1

Contents

Executive Summary

2 Aug 2026
High-risk AI enforcement deadline

€35M / 7%
Maximum fine — prohibited practices

€15M / 3%
Maximum fine — documentation failures

10 Years
Technical File retention (Article 18)

The EU AI Act entered into force on 2 August 2024. For providers of high-risk AI systems, the critical enforcement deadline falls on **2 August 2026** — the date from which Chapter III obligations become enforceable by national market surveillance authorities. Providers whose systems are not supported by compliant Technical Files, completed conformity assessments, and operational post-market monitoring programmes face fines that dwarf the cost of the compliance investment.

This white paper provides a practitioner-level framework for the three obligations that define provider compliance: the Annex IV Technical File (Article 11), the conformity assessment procedure (Articles 43–46), and the post-market monitoring system (Article 72). The central finding is consistent across compliance programmes in every sector we have worked with: organisations that treat these obligations as an integrated programme — building the Technical File contemporaneously, using it as the foundation for conformity assessment, and extending it into a live monitoring programme — achieve compliance at roughly one-fifth the cost of those that address each obligation sequentially and reactively.

This paper is part of a three-paper series. **Paper 2** covers AI governance frameworks, ISO 42001 integration, and the ROI of early governance investment (*unorma.com/ai-governance-framework*). **Paper 3** addresses deployer obligations under Articles 26, 14, 27, and 4

(*unorma.com/deployer-obligations*).

1 The High-Risk Classification Framework

Article 6 and Annex III define eight categories of high-risk AI systems. Classification is determined by the system's **intended purpose** — not its technical architecture or sophistication. A simple linear model used for a purpose falling within Annex III is fully subject to Chapter III; a complex large language model used for general-purpose productivity assistance that does not fall within Annex III is not.

Annex III Category	Scope & Examples	Primary Compliance Trigger
§1 Biometric identification	Real-time remote biometric ID in public spaces; emotion recognition; categorisation by biometric data	Article 5 prohibition check first — if the specific use is prohibited, compliance is not achievable
§2 Critical infrastructure	AI managing electricity grids, water systems, digital infrastructure, road traffic	Safety-critical context triggers maximum Technical File depth and human oversight requirements
§3 Education & training	AI scoring educational assessments, selecting students, monitoring learners during exams	Access to education is a fundamental right; FRIA required for public sector deployers
§4 Employment & workers	CV screening, workforce performance monitoring, task allocation AI, termination decision support	Largest private-sector affected population in EU; widespread FRIA obligation for large employers
§5 Access to essential services	Credit scoring, insurance underwriting, benefit eligibility, social services triage AI	Financial sector expected to face earliest enforcement scrutiny; FRIA mandatory for deployers
§6 Law enforcement	Predictive policing, evidence evaluation, crime risk assessment, polygraph AI	Restricted to law enforcement bodies; additional safeguards and democratic oversight apply
§7 Migration & borders	Asylum application risk scoring, border crossing risk assessment, visa processing AI	Highest fundamental rights exposure; FRIA mandatory; strict data governance requirements
§8 Justice & democracy	AI assisting judicial decisions, democratic process management, election-related AI	Judicial independence implications; strictest human oversight and explainability requirements

Table 1 · Annex III high-risk AI categories with scope, examples and primary compliance trigger.

SCOPE DISCIPLINE

A provider who allows a deployer to use the system beyond its documented intended purpose — without updating the Technical File — becomes liable for all consequences of that expanded use under Article 25. Deployer use-case expansion without provider notification is simultaneously a provider documentation failure and a deployer obligation breach. Scope is both a product decision and a legal boundary.

2 Building the Annex IV Technical File

The Technical File is the central compliance artefact. It must exist before market placement, be maintained throughout the system's operational life, and be retained for ten years after withdrawal. Article 18 makes it available to market surveillance authorities on demand — in practice it is the first document examined in any enforcement investigation.

§	Section Title	Required Content	Automation Potential	Update Trigger
§ 1	General Description	Intended purpose, negative scope, version, affected populations, foreseeable misuse	Low — legal authorship	Any change to intended use or scope
§ 2	Design & Development	Architecture diagrams, compute resources, third-party components, design rationale	Medium — extract from IaC/CI	Architecture or component change
§ 3	Training Data	Dataset provenance, preprocessing, demographic composition, bias evaluation results	High — DVC, dataset registry, Fairlearn	Dataset version update
§ 4	Performance & Monitoring	Model card, pre-determined thresholds, known limitations, logging capabilities	High — MLflow, Weights & Biases	Model retrain with metric change
§ 5	Cybersecurity	AI-specific threat model, adversarial robustness testing, OWASP LLM Top 10 assessment	Medium — ingest pen-test outputs	Architecture or component change
§ 6	Testing Results	Full evaluation results, pre-dated acceptance thresholds, subgroup performance data	High — export from eval pipeline	New evaluation run
§ 7	Deployer Instructions	Post-market monitoring plan, oversight requirements, logging specs, usage conditions	High — observability config export	Operational change
§ 8	Declaration of Conformity	Signed declaration: system identity, standards applied, authorised signatory	None — legal instrument	New conformity assessment

Table 2 · Annex IV sections: required content, automation potential and update triggers.

The Documentation Trigger System

A Technical File that does not update when the underlying system changes is non-compliant regardless of how complete it was at creation. Embed change-triggered documentation updates in your CI/CD pipeline:

- **Model retrain (same architecture, same data):** Auto-update §4 and §6 from the new experiment run; human review only if metrics breach documented thresholds.
- **Dataset version change:** Auto-update §3 provenance records; re-run bias evaluation pipeline; compliance sign-off on fairness output before release.
- **Architecture or component change:** Re-extract §2; flag §5 for adversarial retest; compliance team assesses whether a new conformity assessment is triggered.
- **Intended-use expansion:** Block deployment; update §1; full conformity assessment required before the new deployment context goes live.
- **Upstream GPAI model update:** Alert compliance team; re-run baseline performance and bias tests; assess whether behaviour change constitutes material non-conformity.

TECHNICAL GUIDANCE — PRE-DATING THRESHOLDS

Pre-register acceptance thresholds in your experiment tracking platform **before** evaluation runs execute. In MLflow, log them as run parameters; in Weights & Biases, include them in the run config. This creates the timestamped audit trail demonstrating thresholds preceded results — a point assessors specifically verify, because post-hoc threshold-setting is a well-known data integrity failure mode that immediately raises a conformity assessment red flag.

3 Quality Management System — Article 17

Article 17 requires providers to establish a documented Quality Management System before their first high-risk AI system is placed on the market. The QMS is the governance infrastructure that makes the Technical File credible — assessors look for evidence that the processes documented in the QMS actually produced the artefacts claimed in the Technical File.

QMS Element	Article 17 Requirement	Common Gap Found in Assessments
Regulatory compliance strategy	Documented approach to identifying and meeting current and evolving regulatory requirements	Strategy exists but is not linked to specific Articles — appears generic rather than AI Act-specific
Design & development controls	Systematic procedures for design decisions including documentation requirements at each stage	Procedures documented but not followed — no evidence of procedure-governed decisions in Technical File
Data governance procedures	Controls for training, testing and validation data quality, representativeness and bias management	High-level policy exists; no evidence of per-dataset governance records required by Article 10
Human oversight design	Process for designing and validating the five Article 14(4) oversight capabilities into products	Oversight described as a deployment consideration; no design-stage engineering controls documented
Post-market monitoring system	Systematic process for collecting and analysing production performance data per Article 72	Monitoring plan listed in Technical File §7 but infrastructure not built at conformity assessment time
Incident management	Procedures for capturing, classifying, and reporting serious incidents per Article 73 timelines	Incident procedures exist for general software; AI-specific 15/3-day reporting workflow not defined
Corrective action	Process for implementing and documenting corrective actions with Technical File linkage	Corrective actions implemented but not linked back to Technical File — creates divergence between documentation and live system

Table 3 · Article 17 QMS elements with common assessment findings.

4 Conformity Assessment Pathways

Articles 43–46 define two conformity assessment pathways. The choice depends on the system's Annex III category — it is not a free choice based on provider preference or resource constraints.

ay	When Required	Key Steps	Timeline
ment (Annex	Most Annex III categories — the default for employment AI, essential services AI, education AI, justice AI	1. Establish QMS per Art.17 2. Complete Annex IV Technical File 3. Conduct Annex VI assessment — review Technical File against Articles 8–15 4. Sign Declaration of Conformity (Annex V) 5. Register in EU AI database	3–6 months from
Body VII)	Real-time biometric ID systems; law enforcement AI requiring NB under Art.43(1); any system where provider elects external validation	1. Select EU-accredited Notified Body 2. Submit Technical File 3. Technical examination of system 4. QMS audit against Article 17 5. Conformity certificate issued 6. Declaration + Registration	6–12 months; NB is a constraint
er-Level systems)	Deployers using GPAI-based systems where provider has satisfied Article 53 GPAI transparency obligations	1. Verify provider GPAI compliance documentation 2. Conduct deployer-specific conformity assessment for system-level requirements 3. FRIA per Article 27 where deployer category triggers it	1–3 months if pro docs complete

Table 4 · Conformity assessment pathways by system type and Annex III category.

NOTIFIED BODY AVAILABILITY WARNING

Notified Body queues for Annex VII assessments are congested in several EU member states as of Q1 2026. Providers who require Notified Body assessment and have not already submitted applications face a structural risk of missing the August 2026 deadline regardless of how well prepared their Technical File is. Self-assessment pathways are available for most systems — but require a complete, operational QMS which itself takes 8–12 weeks to establish from scratch.

5 Post-Market Monitoring — Articles 72 & 73

Article 72 requires providers to maintain a post-market monitoring system from the day the system enters service. The system must actively and systematically collect data to evaluate whether the AI system continues to comply with Chapter III requirements throughout its lifetime. A monitoring plan that exists as a document but has no operational infrastructure behind it does not satisfy Article 72.

The Three Drift Types Your Monitoring Must Cover

ype	Definition	Detection Method	Compliance Risk
ift ate Shift)	Statistical distribution of input features shifts away from training distribution	Population Stability Index, KS test, Jensen-Shannon divergence on input features vs. training baseline — tools: Evidently AI, whylogs	Performance claims in Technical File no longer hold for the current populati
t Drift	The underlying relationship between inputs and outputs changes in the real world — the model's learned function becomes incorrect	Rolling prediction error rates where ground truth available; confidence distribution shifts; operator override rate trends where ground truth delayed	Model predictions become systematic incorrect — Article 9 risk managemen if drift is not detected and corrected
raphic Drift	Disparate impact metrics shift across protected groups even when aggregate performance appears stable	Continuous monitoring of disaggregated fairness metrics — disparate impact ratio, equalised odds differential — against Technical File baselines	Article 10 data governance non-compl discrimination risk; corrective action re before next deployment cycle

Table 5 · Model drift types with detection methods and compliance risk implications.

Article 73: Serious Incident Reporting Obligations

Incident Category	Article 73 Classification	Reporting Deadline	Report Recipient
Death or life-threatening injury directly attributable to AI system output	Serious incident — life-threatening	3 calendar days from awareness	National Market Surveillance Authority
Serious health impairment, property damage, or environmental harm attributable to AI system	Serious incident — non-life-threatening	15 calendar days from awareness	National Market Surveillance Authority
Serious disruption of critical infrastructure operations	Serious incident — infrastructure	15 calendar days from awareness	National MSA + relevant sector regulator
Performance degradation triggering Technical File non-compliance without serious harm	Compliance malfunction — internal only	No external reporting; corrective action + Technical File update + deployer notification required	Internal — deployers notified

Table 6 · Article 73 serious incident classification, reporting deadlines and report recipients.

MONITORING INFRASTRUCTURE MUST PRECEDE LAUNCH

Article 72's requirement for an active and systematic monitoring system means the infrastructure — dashboards, alert thresholds, incident capture workflow, corrective action governance — must be operational on the day the system goes live. A monitoring plan written after an incident is not a monitoring plan; it is an incident report. Build monitoring infrastructure in parallel with the system, not after deployment.

6 The 2026 Compliance Roadmap

Phase	Period	Key Activities	Gate Condition
Phase 1 Foundation	Now – Apr 2026	Complete AI system inventory and Annex III classification. Assign compliance owners for each high-risk system. Establish QMS skeleton. Begin dataset versioning (DVC or equivalent) and model registry integration. Identify Notified Body requirement.	Every high-risk system has a named owner and draft scope statement
Phase 2 Documentation	Apr – Jun 2026	Complete Annex IV Technical File for each high-risk system. Run mock audit simulation to identify conformity gaps. Remediate critical findings. Pre-register performance thresholds. Conduct bias evaluation and document results.	Technical File passes internal mock audit with no critical findings
Phase 3 Assessment	Jun – Jul 2026	Conduct internal conformity assessment (Annex VI) or submit to Notified Body. Complete FRIA for qualifying deployers. Sign Declaration of Conformity. Register in EU AI database.	Declaration of Conformity signed and EU database registration confirmed
Phase 4 Operational	Aug 2026 →	Activate post-market monitoring programme. Train oversight personnel to Article 14 standard. Integrate incident reporting workflow with Article 73 timelines. Conduct quarterly compliance reviews and annual Technical File audits.	Active monitoring data flowing; training records complete for all oversight personnel

Table 7 · Phased compliance roadmap to the August 2026 enforcement deadline.

Conclusion

EU AI Act compliance for high-risk AI providers is not achievable in weeks. The organisations that will clear the August 2026 deadline with confidence are those that started building documentation infrastructure in 2025 — treating compliance as a product requirement rather than a pre-launch legal exercise. The three obligations in this paper are most efficiently addressed as an integrated programme: the dataset registry that feeds §3 also feeds the monitoring baseline; the evaluation pipeline that produces §6 also produces the post-market monitoring reference; the QMS that governs the conformity assessment also governs the corrective action process.

About Unorma

Unorma is the compliance infrastructure platform for high-risk AI providers. Our Document Generator automates Annex IV Technical File production from your existing ML infrastructure — connecting to MLflow, DVC, Fairlearn, and your CI/CD pipeline to auto-populate all eight Annex IV sections. Our Audit Simulation identifies non-conformity findings before your formal assessment. Our AI System Inventory maintains your complete compliance record in an immutable evidence vault accessible to market surveillance authorities within 48 hours.

→ Technical File automation: unorma.com/features/document-generator

→ Audit simulation: unorma.com/features/audit-simulation

→ Full white paper series: unorma.com/resources

[1] EU AI Act penalty structure: Article 99. Tier 1 (prohibited practices): €35M or 7% global turnover. Tier 2 (other violations): €15M or 3%. Tier 3 (false information): €7.5M or 1%.

[2] Notified Body availability data: CEN-CENELEC JTC 21 working group updates and industry surveys, Q1 2026.

[3] Compliance cost estimates: industry benchmarks adjusted from ISO 27001 and GDPR implementation programmes.