unorma

WHITE PAPER 3 OF 3

Deployer Obligations: Human Oversight, FRIA & AI Literacy

A Practical Guide to Articles 26, 14, 27, and 4 for EU AI Act Deployers in 2026

3

March 2026 · unorma.com/deployer-obligations

Contents

Executive Summary

Art. 26
9 distinct deployer obligations from day of deployment

Art. 14
5 enumerated oversight capabilities — must be product-designed

Art. 27
FRIA mandatory: public bodies, banks, insurers, large employers

Art. 4
AI literacy: role-specific, system-specific, documented

The EU AI Act's deployer obligations are systematically misunderstood. The most common misconception: purchasing an AI system from a compliant provider transfers the compliance obligation to the provider. It does not. Article 26 establishes a distinct deployer obligation set that applies from the moment a high-risk AI system is put into service — independent of the provider's compliance status.

This paper addresses three deployer obligations that are simultaneously the most operationally demanding and the most commonly addressed inadequately: human oversight under Article 14, Fundamental Rights Impact Assessments under Article 27, and AI literacy training under Article 4. The central finding: effective deployer compliance is a product design and organisational capability challenge, not a documentation exercise. Oversight that is nominally in place but not embedded in product design is not compliant oversight. A FRIA completed after deployment is not a compliant FRIA.

Paper 1 covers provider obligations — Technical File, conformity assessment, post-market monitoring (*unorma.com/eu-ai-act-compliance*). **Paper 2** covers AI governance frameworks, ISO 42001, and the ROI of early compliance investment (*unorma.com/ai-governance-framework*).

1 Article 26: The Deployer's Nine Obligations

Article 26 defines the deployer's primary obligation set. These nine obligations apply to any deployer of a high-risk AI system — regardless of whether the system was purchased from a provider, built internally, or delivered through a SaaS platform.

Art. 26	Obligation	What It Requires in Practice	Most Common Failure Mode
26(1)	Use within intended purpose	Deploy only for use cases documented in the provider's Technical File; notify provider before any use-case expansion	Use-case scope creep in production — the system is applied to decisions not covered by the Technical File, without provider notification
26(2)	Human oversight measures	Assign technically qualified oversight persons; implement the oversight controls the provider's system design supports; verify operators can exercise genuine evaluation	Oversight assigned nominally to an already-overloaded team with no training on the specific system, and an interface that doesn't support genuine output evaluation
26(3)	Monitor performance	Actively monitor system performance in production; report relevant information and anomalies to the provider; maintain deployer-side incident capture	Relying entirely on provider monitoring; no deployer-side operational visibility; incidents captured only after they cause harm
26(5)	Log retention	Retain automatically-generated operational logs for minimum 6 months (Article 12); ensure AI decision logs are excluded from general log rotation policies	AI decision logs deleted after 30-day general log rotation; provider's Technical File §7 logging specification not implemented in deployer infrastructure
26(6)	DPIA and FRIA	Conduct GDPR DPIA per Article 35 where personal data processed; conduct FRIA per Article 27 where deployer category triggers it — public bodies, financial services, large employers	Assuming provider's FRIA covers deployer's specific deployment context; not conducting deployer-specific assessment
26(7)	Inform individuals	Inform natural persons subject to AI-assisted decisions that they are interacting with or being assessed by an AI system	Assuming provider's general product disclosures satisfy deployer's specific individual notification obligation
26(8)	Deployer instructions compliance	Implement the human oversight, usage conditions, and operational requirements set out in the provider's deployer instructions (Technical File §7)	Deployer instructions reviewed by legal at procurement but never operationally implemented in the deployment environment
26(9)	Notify supervisory authority	Notify the relevant market surveillance authority when the system creates unacceptable fundamental rights risk in the specific deployment context	No defined escalation path to the supervisory authority; no clarity on which authority has jurisdiction for the specific deployment

Table 1 · Article 26 deployer obligations: requirements and common failure modes.

2 Article 14: Engineering Meaningful Human Oversight

Article 14 is satisfied when the product is designed so that human oversight is technically meaningful — not when a policy document states that oversight occurs. The five enumerated capabilities in Article 14(4) define the minimum technical controls that must be present in the product interface and training programme for oversight to meet the legal standard.

Art. 14(4)	Legal Requirement	Engineering Implementation	Evidence for Assessors
(a) Capabilities & limitations	Oversight persons must genuinely understand what the system can and cannot do before reviewing its outputs in any given context	Decision context panel in the operator UI showing system confidence, relevant feature inputs, and documented limitations for the decision type; capability/limitation summary accessible within one click of every review screen	UI screenshot in Technical File §7; capability description in deployer instructions; training records confirming system-specific training
(b) Anomaly monitoring	Operators must have technical means to detect when the system behaves unexpectedly for the current input	Anomaly indicator in operator interface triggered when input falls outside training distribution; confidence score display with calibration context; distribution shift alerts from monitoring pipeline piped to operator dashboard	Monitoring dashboard specification; alert threshold documentation; anomaly event log showing operational monitoring is active
(c) Automation bias awareness	Operators must be specifically informed about automation bias and trained to actively resist it — not just generally aware of AI limitations	Training module with scenario-based automation bias exercises (not just conceptual content); UI friction for high-confidence outputs on high-stakes decisions (structured review question before acceptance)	Training curriculum documentation; completion records per operator per system; UI friction specification
(d) Output interpretation	Operators must understand what the AI's output means — what confidence scores indicate, what factors drove the decision, and appropriate use contexts	Explainability features in operator interface; calibration-aware confidence display; per-decision feature importance; output interpretation guide accessible within the UI	Explainability method documented in Technical File §4; interface specification; output interpretation guide content
(e) Decide not to use output	Operators must be technically and procedurally able to disregard, modify, or override AI output and have the override recorded with attribution	Override workflow with equal visual prominence to acceptance (not a buried "reject" option); structured rationale capture; divergence logging with final human decision attribution and timestamp	Override log records demonstrating overrides actually occur; override workflow UI specification; divergence rate in post-market monitoring data — absence of overrides is a compliance red flag

Table 2 · Article 14(4) capabilities: legal requirements, engineering implementations, and assessor evidence.

THE OPERATIONAL REALISM TEST

For every Article 14 capability, apply this test: "If a market surveillance authority came tomorrow and asked to see this capability in operation, could we demonstrate it?" For capability (e), this means producing actual override log records — not a policy. An absence of override entries across thousands of decisions is a significant red flag suggesting operators are accepting AI outputs without genuine evaluation, which is the compliance failure Article 14 is designed to prevent.

3 Automation Bias: The Compliance Risk Inside Every Oversight Process

Automation bias — the tendency to over-rely on automated outputs and under-apply independent judgment — is the mechanism through which nominally compliant human oversight fails in practice. Article 14(4)(c) specifically requires that oversight persons be made aware of automation bias, making it the only cognitive bias explicitly named in the legislation.

The research on automation bias mitigation is unambiguous on one point: **conceptual training alone is insufficient**. A training module that explains what automation bias is and why it occurs produces minimal durable behaviour change. Training that includes scenario-based exercises — where operators practice detecting and overriding AI errors in realistic cases — produces substantially better 90-day retention and real-world behaviour change.

Design Pattern	How It Works	Evidence of Effectiveness
Scenario-based training with incorrect AI outputs	Training exercises present operators with AI outputs that are confidently wrong — cases that would produce automation bias in non-trained operators — and require learners to identify the error and document the override rationale	Ergonomics journal research (2022) found scenario-trained operators showed durable reduction in automation bias at 90-day follow-up vs. minimal retention in conceptual-only training group
UI friction for high-stakes decisions	When the AI's confidence score exceeds a threshold for decisions with high consequences, the interface inserts a structured review question before acceptance — requiring operators to articulate why they agree with the AI output before they can proceed	Google PAIR research on human-AI interaction shows structured intervention at acceptance significantly reduces automatic acceptance rates for high-confidence outputs
Confidence score calibration display	Rather than showing raw confidence percentages (which operators often misinterpret as accuracy), display calibration context: "This system is correct 87% of the time when it reports this confidence level for this decision type"	Nature Digital Medicine study on AI-assisted radiology found miscalibrated confidence displays directly correlated with over-reliance on incorrect outputs
Override rate monitoring as a compliance signal	Track override rates by operator, decision type, and confidence level. Falling override rates over time — particularly for high-confidence outputs — signal increasing automation bias and should trigger retraining before they become a compliance finding	Acts as a leading indicator of Article 14 compliance degradation: override rate trends feed directly into post-market monitoring data for Art. 72 reporting

Table 3 · Automation bias mitigation design patterns with evidence of effectiveness.

4 Article 27: Fundamental Rights Impact Assessment

Article 27 requires specific categories of deployers to conduct a Fundamental Rights Impact Assessment (FRIA) before deployment. The obligation extends more broadly into the private sector than most organisations initially assume.

Deployer Category	FRIA Required?	Sector Examples	Scope Note
Public bodies & public authorities	Yes — always, for any high-risk AI use	Government agencies, municipalities, tax authorities, courts, law enforcement, social services	Any use in the exercise of public powers or provision of public services
Private entities — essential services	Yes — for Annex III §5 systems	Banks, insurers, credit institutions, healthcare providers, energy operators, transport operators	AI used for credit scoring, insurance underwriting, healthcare triage, benefit eligibility decisions
Large-scale private employers	Yes — where systematic and at scale	Large companies using AI for recruitment, performance management, workforce analytics	"Large-scale" threshold not defined — document your classification reasoning explicitly
All other private deployers	Not mandatory — strongly advised	All other private sector companies using high-risk AI systems	Proactive FRIA demonstrates due diligence; increasingly required in enterprise B2B procurement

Table 4 · Article 27 FRIA obligation by deployer category.

EU Charter Rights Most Commonly Implicated by High-Risk AI

Charter Right	Art.	Typical AI Risk Scenario	FRIA Assessment Focus
Human Dignity	1	Dehumanising treatment in automated systems; systems treating individuals as data points	Does the system's design and deployment context respect the inherent worth of affected individuals?
Non-Discrimination & Equality	21–23	Disparate impact on protected groups in hiring, credit, housing, education, or justice decisions	Are there demographic disparities in AI outputs? Are mitigations proportionate to the identified risk?
Privacy & Data Protection	7–8	Excessive data collection; use beyond original purpose; re-identification risk from AI inference	Does the system process personal data proportionately and with appropriate technical safeguards?
Access to Justice & Effective Remedy	47	AI decisions that lack explainability or that individuals cannot challenge through accessible processes	Can affected individuals understand and effectively challenge AI-assisted decisions?
Right to Work & Fair Employment	15, 31	AI-driven hiring and termination decisions; automated performance management affecting working conditions	Are AI employment decisions fair, transparent, explainable, and contestable by affected workers?
Rights of the Child	24	AI in education, content recommendation for minors, parental decision support AI	Are the best interests of the child treated as a primary consideration in the system's design and deployment?

Table 5 · EU Charter rights most commonly implicated by high-risk AI, with FRIA assessment focus.

5 Combined FRIA + DPIA: Section-by-Section Template

Article 27(5) explicitly permits the FRIA and GDPR DPIA to be conducted as a single combined assessment. For AI systems processing personal data — which is nearly all high-risk AI — the combined approach is recommended: the FRIA structure (broader scope) leads, with the GDPR-specific data protection risk assessment embedded as a sub-section.

Section	Content Required	GDPR DPIA Integration
S1 — Identification	Deployer legal form and Art.27 category; AI system identity and provider reference; EU database registration; assessment date, lead assessor, approval authority	Add: DPO identity; GDPR lawful basis; data controller/processor status
S2 — Process Description	Process supported by AI; decision types and AI's role; deployment frequency and volume; geographic scope; deployment duration	Add: categories of personal data processed; data flows; data retention periods; international transfers
S3 — Affected Population	Categories of persons subject to AI-assisted decisions; specific attention to vulnerable groups; for each group: nature of decision, potential consequences, and particular vulnerability to AI-related harm	Add: data subject categories under GDPR; special categories identification; children affected
S4 — Rights Risk Assessment	For each relevant Charter right: risk scenario, likelihood, severity, and mitigation status. Minimum: human dignity, non-discrimination, privacy, access to justice, right to work where applicable, rights of the child where applicable	Embed GDPR Art. 35 risk assessment here: risks to data subject rights under Arts. 5-9, Art. 22 automated decision-making rights

Section	Content Required	GDPR DPIA Integration
S5 — Safeguards	For each identified risk: specific safeguard; technical/operational/legal classification; named responsible person; human oversight and individual redress mechanisms; monitoring commitments	Add: GDPR Art. 22 right to human review mechanism; Arts. 15-22 data subject rights exercise procedures; DPO consultation record
S6 — Conclusion	Residual fundamental rights risk after mitigations; deployment decision; prior consultation commitment where significant residual risk; review schedule	Add: supervisory authority prior consultation under GDPR Art. 36 where required; DPO sign-off

Table 6 · Combined FRIA + DPIA template: section structure with Article 27 and GDPR Art. 35 content requirements.

6 Article 4: AI Literacy — Role-Differentiated Framework

Article 4 requires deployers to take measures ensuring staff and other persons dealing with the AI system have sufficient AI literacy — appropriate to their role, the system's technical complexity, and the deployment context. For high-risk AI oversight personnel, Article 14(4) creates a substantially higher competency bar.

Target Roles	Core Competencies	Format	Duration	Refresh Trig
All staff whose work may be affected by or involves use of AI tools — organisational baseline	What AI is and is not; how AI decisions differ from human decisions; AI limitations and error patterns; personal rights when subject to AI; how to escalate concerns	Self-paced e-learning + knowledge check	1–2 hrs	Annual or on new AI syste
Employees actively using AI systems in workflow — including managers receiving AI recommendations	All Tier 1 + this system's documented capabilities and limitations; failure mode awareness; automation bias and mitigation; confidence score interpretation; override procedure; documentation requirements	Blended: foundation e-module + practical workshop with sandbox system including deliberately wrong AI outputs to detect	4–6 hrs per system	On system d version upda
Designated Art. 14 oversight persons; AI governance leads; compliance and legal staff; product owners and ML engineers for high-risk systems	All Tier 1–2 + Art. 14(4) capabilities in operational depth; monitoring dashboard interpretation; incident classification and Art. 73 reporting workflow; Technical File threshold documentation; FRIA methodology; corrective action governance; penalty structure and liability framework	Instructor-led with live system exercises, monitoring dashboard walkthroughs, mock incident classification, and compliance scenario workshops	16–24 hrs per role type	On role assig system versi annually

Table 7 · Three-tier AI literacy framework: role targeting, competencies, delivery format, duration, and refresh triggers.

Training Documentation Requirements

Training records must be linked to the specific AI system **and version** for which they were completed. For each oversight person and each system, the compliance record should specify: system name and version; training module name and version; completion date; assessment result; and next refresh date. Records must be stored in the compliance evidence vault and producible within 48 hours on regulatory request.

CONTRACTOR AND THIRD-PARTY STAFF

Article 4 covers "staff and other persons dealing with the operation of AI systems on their behalf" — explicitly including contractors, outsourced service providers, and third-party operators. Training obligations follow operational responsibility, not employment status. Include AI literacy training requirements in contracts with service providers who operate your high-risk AI systems; require completion records to be provided to you for your compliance documentation. A contractor operating your high-risk AI system without adequate training is a compliance risk you are responsible for.

Conclusion

Deployer compliance under the EU AI Act is not satisfied by policy documents and generic training programmes. It is satisfied when product design embeds the five Article 14(4) oversight capabilities as genuine technical controls; when the FRIA is conducted before deployment with authentic assessment of Charter rights risk; and when AI literacy training is role-specific, system-specific, and documented with completion records linked to each system version.

The deployers that will demonstrate effective compliance under market surveillance scrutiny are those that have built these requirements into their deployment processes — producing the override logs, training records, and FRIA documentation that regulators will examine. The organisations that have not will face a combination of reactive remediation costs, enforcement exposure, and reputational consequences that dwarf the cost of the proactive investment described in this paper.

About Unorma

Unorma's compliance platform supports deployers with structured FRIA workflows aligned to Article 27, training record management linked to your AI system inventory, Article 14 compliance assessment for your product's oversight design, and an evidence vault storing all deployer compliance records in tamper-evident, regulator-accessible storage.

→ FRIA & deployer compliance tools: unorma.com/features/document-generator

→ AI System Inventory: unorma.com/features/ai-system-inventory

→ WP1 — Provider obligations: unorma.com/eu-ai-act-compliance

→ WP2 — Governance & ISO 42001: unorma.com/ai-governance-framework

[1] EU AI Act Article 26 deployer obligations; Article 14 human oversight; Article 27 FRIA; Article 4 AI literacy.

[2] Automation bias research: Parasuraman & Riley (1997); Cummings (2004); Nature Digital Medicine AI radiology study (2023); Ergonomics journal scenario training study (2022).

[3] FRIA/DPIA integration: Article 27(5) and Recital 89, EU AI Act. EDPB guidelines on data subject rights, 2022.