ANNEX TO INDEPENDENT RESEARCH

# The Deployer Multiplier

## Three downstream segments discovering Article 26 — HR, healthcare, retail banking

The parent *EU AI Act Readiness Report* (Sprinkling Act, April 2026) screened 50 European AI providers and identified that the cascade effect from a small number of high-risk providers downstream to thousands of European deployer organizations is "often the most actionable finding of this analysis" (parent §3.1). This annex develops that cascade for three downstream segments (human-resources deployers, healthcare deployers, and retail-banking deployers) that the parent identified as exposed but did not analyse from the deployer's seat.

| 3 | 3 500+ | 12 | 4 |
|---|---|---|---|
| Deployer segments | EU deployers exposed (parent §3.1) | Art. 26 paragraphs mapped | Path-to-discovery triggers |

| Annex Details | |
|---|---|
| **Date prepared** | May 6, 2026 |
| **Drafting period** | April 15 – May 6, 2026 |
| **Parent document** | EU AI Act Readiness Report (April 2026) |
| **Parent sample** | 50 companies (45 valid) |
| **Parent sectors** | 4 — HRTech, MedDev, FinTech/InsurTech, Public AI |
| **Parent countries** | 15+ — FR, DE, BE, NL, PL, DK, CH, IE, UK |
| **Annex scope** | Art. 25, 26, 27 deployer cascade |
| **Annex segments** | HR, Healthcare, Retail Banking |
| **Analyst** | Lamar B. Shucrani, Founder |
| **Independence** | No commercial relationship |

# 0. Scope notes and methodological refinements

*This annex builds on the parent* EU AI Act Readiness Report *(Sprinkling Act, April 2026). The notes below clarify scope, methodology, and the relationship between the parent's findings and this annex's analysis. They do not amend the parent report's substantive findings.*

### 0.1 Pivot from provider screening to deployer cascade

The parent report screens fifty European AI *providers* against six regulatory gates (G1–G6). The parent's §3.1, "The Deployer Multiplier," identifies that each high-risk provider serves downstream deployers who inherit Article 26 obligations, and notes that the cascade is "often the most actionable finding of this analysis." This annex develops that cascade. Its perspective is the European deployer's, not the provider's. The five sectoral cases in the parent report (§4) are not re-analysed here; instead, three downstream deployer segments are analysed.

### 0.2 Deployer scope under the AI Act

Article 3(4) defines a "deployer" as any natural or legal person under whose authority an AI system is used in the course of professional activity. The household exemption is narrow. Most European organizations using purchased AI tools (credit-scoring engines, scribe applications, screening platforms) meet the deployer definition without recognising it. This annex analyses three segments where that recognition gap is most visible.

### 0.3 Segment selection criteria

Three segments are analysed: human-resources deployers (HR directors and CHROs), healthcare deployers (hospitals and clinics using AI scribe applications qualified as Medical Device Software under MDR/IVDR), and retail-banking deployers (credit institutions using third-party scoring). Each meets three cumulative criteria: (a) the segment maps to a named high-risk classification under the AI Act (Annex III §4, Article 6(1) MDR/IVDR pathway, and Annex III §5(b) respectively); (b) European data on segment size is available, directly via supervisory registers or via parent-report extrapolation; (c) the segment's path-to-discovery is non-trivial — ordinary operational practice does not surface the deployer obligation without an external trigger.

### 0.4 Anonymisation and case methodology

Where European public registers permit identification of deployer entities (banking via ACPR/EBA registers; healthcare via national hospital and CPAM lists), this annex constructs anonymised composite profiles in the manner of the parent report's §4 — descriptors drawn from publicly available information, no specific organization named. Where public registers of deployer entities do not exist (HR-AI deployers), the relevant segment is illustrated typologically; the methodological dependency is stated explicitly in §3 below.

# 0. Scope notes (continued)

### 0.5 Reconciliation note on the parent's Article 26 obligations list

The parent *EU AI Act Readiness Report* §3.1 enumerates the obligations that downstream deployers inherit and includes the phrase "fundamental rights impact assessment where applicable" within that list. The FRIA obligation is, in the regulation as enacted, the subject of Article 27 rather than Article 26. The two articles operate in parallel for the deployer: Article 26 sets the operational duties applicable to all high-risk deployers (with paragraph-specific triggers), and Article 27 sets the FRIA obligation for the specific deployer categories listed in Article 27(1). This annex applies that distinction throughout (§5.5, §6.5 treat FRIA under Article 27); the parent's terminology is preserved as published, and the reconciliation is recorded here for readers reading the two documents in sequence.

### 0.6 Carry-forward of the parent's certification blind-spot finding

The parent §2 Finding 3 reports that, of fourteen entities in the screened sample holding existing regulatory certifications (CE MDR, CE-IVD, banking licence, ISO management certifications), zero had publicly mapped their AI Act position. The same blind spot propagates to the deployer layer analysed here: organizations that have invested heavily in adjacent compliance regimes (GDPR, ISO 27001, MDR/IVDR) frequently assume the AI Act is absorbed by those regimes, and have not produced the deployer-side cascade that the AI Act introduces as a distinct artefact. The parent's "0 of 14" provider-level signal is the upstream half of the same observation this annex develops downstream.

### 0.7 The parent multiplier figure — auditability note

The parent report's §3.1 derives the "more than 3,500 European organizations" figure from publicly stated client counts of four named high-risk providers within the parent's sample of 50 screened companies. The "tens of thousands" extrapolation across the 37 high-risk providers screened is an order-of-magnitude estimate, not a direct measurement; the "hundreds of thousands EU-wide" figure aggregates beyond the parent's sample using publicly disclosed market reach.

This annex adopts those figures verbatim as the parent's published estimates and does not re-derive them. The four providers underlying the 3,500+ figure correspond to entries in the parent's §6 sectoral tables whose European client counts are publicly disclosed on the providers' corporate websites or in their investor materials at the time of the parent's screening (April 2026). The parent report, by editorial choice, does not name the four providers in its published text; the underlying screening data is retained by Sprinkling Act and made available to qualified reviewers (regulators, peer researchers, journalists) on request to **contact@sprinklingact.com** for the purpose of replicating the multiplier derivation. A future revision of this annex (v1.1) will list the four providers by their parent-§6 row identifiers and their respective publicly disclosed counts, conditional on prior re-verification of those counts at time of revision (provider disclosures evolve).

### 0.8 Article 26 paragraphs as cascade, not checklist

Article 26 contains twelve paragraphs that differ in trigger, scope, and enforcement posture. Treating Article 26 as a single checklist item underestimates it by an order of magnitude. The reference table in §3 below consolidates the twelve paragraphs and tags the operational trigger of each. For each of the three deployer segments, the active subset is identified, typically four to seven paragraphs simultaneously, not all twelve.

### 0.9 Annex version note

Annex A v1.0 (May 6, 2026), first edition; drafting window April 15 – May 6, 2026. Develops the parent report's §3.1 ("The Deployer Multiplier") for three downstream segments. Full document metadata in §12.

# 0. Scope notes (continued)

### 0.10 Note on Digital Omnibus negotiations

The EU's *Digital Omnibus on AI* (the Commission's simplification package proposing, among other measures, postponement of certain AI Act deadlines) is in ongoing trilogue between the European Parliament, the Council of the EU (under the Cypriot Presidency until June 30, 2026), and the European Commission. The first political trilogue closed without agreement; a second political trilogue convened on April 28, 2026 also failed to reach agreement after extended negotiations, and a further trilogue has been scheduled for May 13, 2026. The Cypriot Presidency has signalled its intention to close the file before its term ends.

Public reporting from the April 28, 2026 trilogue indicates that the breakdown was not on the most prominent element of the proposal — the postponement of Annex III high-risk standalone deadlines, but on the conformity-assessment architecture for AI systems embedded in products covered by existing EU sectoral safety law (Annex I). The implication for deployers in the three segments analysed in this annex is direct: the postponement of Article 26 deployer obligations and Article 27 FRIA triggers depends on adoption of the Omnibus, which depends in turn on resolution of an Annex I question that is not, in itself, a deployer matter.

If the trilogue does not conclude with formally adopted text published in the Official Journal before August 2, 2026, the originally enacted timeline of Regulation (EU) 2024/1689 takes effect — Article 26 §1–§12, Article 27 (where applicable), Article 49 (EU-database registration for public-law bodies), and Article 50 (Article 50 transparency for deployers) become binding on that date for the systems in scope. This annex therefore treats August 2, 2026 as the binding deadline. Any Omnibus-driven extension is treated as bonus runway, not as an active premise. The status reported here reflects publicly available information as of May 6, 2026 and may evolve before final adoption.

FIGURE 1 — AI ACT DEPLOYER-CASCADE TIMELINE (AS OF MAY 6, 2026)

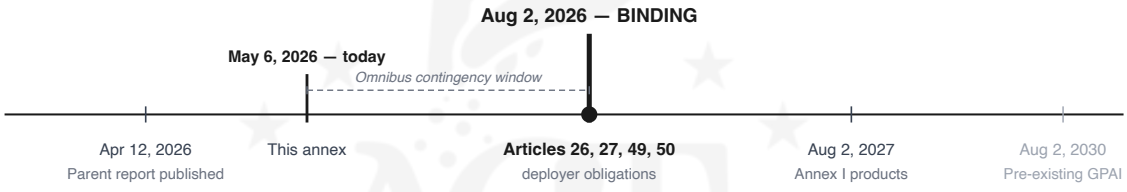


**Apr 28, 2026** · Trilogue 2 closed without agreement **May 13, 2026** · Trilogue 3 scheduled **Jun 30, 2026** · Cypriot Presidency ends

**Reading.** The Omnibus must be formally adopted and published in the Official Journal before August 2, 2026 to defer Article 26, 27, 49 and 50. Three contingency events condition that adoption window. If they do not converge in time, the original AI Act timeline applies as enacted. Sources: A&O Shearman, Modulos, IAPP, Bird & Bird; full citations §10.

### 0.11 What this annex is not

This annex is not legal advice. It is not a conformity assessment under Article 43 AIA. It is not a certification of any kind. It is not a substitute for sector-specific guidance (MDCG for medical devices, EBA for banking, national employment law for HR). Its purpose is to make explicit the cascade that ordinary operational documentation does not currently produce, so that an organization recognising itself in one of the three segments can begin work earlier rather than later. Specific classifications for any specific deployer require a tailored screening.

# 1. Framing — The downstream multiplier

The parent *EU AI Act Readiness Report* identified that the regulatory exposure of European AI providers does not stop at the provider. It cascades. Each high-risk provider serves a downstream population of deployers who, by purchasing the provider's system and using it under their authority, inherit the obligations of Article 26 — risk management measures, human oversight, input-data quality where applicable, ongoing monitoring, automatic logging, workplace-deployment information, registration in the EU database for public-law bodies, integration with the GDPR Article 35 DPIA, post-remote biometric limits where applicable, information of natural persons subject to high-risk decisions, and cooperation with competent authorities.

The parent reports that, from publicly stated client counts of just four named high-risk providers in its sample, more than 3,500 European organizations are already identifiable as deployers exposed to Article 26. Across the 37 high-risk providers screened, the realistic order of magnitude is in the tens of thousands. EU-wide, beyond the parent's sample, the figure runs into the hundreds of thousands. *Most of these organizations do not know they qualify as deployers under the AI Act.* This observation is consistent with industry readiness surveys published in 2025–2026 by professional associations and consultancies active in EU AI compliance, which converge on the finding that a clear majority of organizations using third-party AI tools have not formally evaluated their AI Act deployer status; precise survey methodologies and population coverage vary, and the surveys are not pooled here as a single statistic.

### 1.1 Silence is not absence

An organization's silence on its deployer status is not evidence that it has internally analysed the question and concluded that it does not apply. In the great majority of cases, the question has not been posed. The organization purchases an AI tool, integrates it into a workflow, and treats compliance as the provider's problem. Article 26 is structured against precisely this assumption. The provider's CE marking, the provider's technical file, the provider's instructions for use are inputs to the deployer's obligations, not substitutes for them.

### 1.2 Three orders of magnitude

The cascade is not abstract. The parent report's three orders of magnitude. four providers → 3,500 deployers; thirty-seven providers → tens of thousands; EU-wide → hundreds of thousands — describe a population that, on August 2, 2026, will become subject to a deployer obligation the largest fraction of which has never been formally examined for them. The "EU-wide → hundreds of thousands" extrapolation inherits the parent report's sample composition (approximately seventy percent French; parent §9.4) and scales the parent's provider-side observations to the broader Union deployer population using publicly disclosed market reach. The extrapolation is order-of-magnitude rather than a point estimate; a reader weighting non-French member states differently would arrive at a different point estimate but, on plausible reweightings, not at a different order of magnitude. The population is the structural problem.

### 1.3 Three business consequences for deployers

The same three consequences identified in the parent report for providers cascade onto the deployer, with operational adaptations. **Contract risk.** European business-to-business clauses increasingly require an article-mapped AI Act position on demand; a deployer that has not produced a deployer-side position cannot meet that contractual deliverable, regardless of how robust the provider's documentation is. **Supervisor exposure.** In regulated sectors (banking under ACPR/EBA, healthcare under national authorities and Notified Bodies, employment under labour inspectorates and the Court of Justice's expanding case law on automated decision-making) the absence of a documented deployer position is itself a finding when the supervisor inspects AI controls. **Deployment blocker.** Article 26 §7 (works-council consultation), §8 (EU-database registration for public bodies), and Article 27 §3 (pre-deployment FRIA notification) are gating events. They cannot be retroactively documented after a system is in use; they have to occur before deployment or before next material change.

# 1. Framing (continued)

### 1.4 The three segments analysed in this annex

Three segments are analysed in turn. **Segment 1 (§4)**: human-resources deployers (HR directors and CHROs of medium to large European employers using AI tools in recruitment, performance evaluation, task allocation, or workforce monitoring. The segment maps to Annex III §4. **Segment 2 (§5)**: healthcare deployers) public hospitals, conventioned private clinics, and commercial private practices using AI scribe applications, diagnostic decision support, or clinical-documentation generation qualified as Medical Device Software under MDR/IVDR. The segment maps to Article 6(1) AIA via the MDR/IVDR safety-component pathway; it is therefore high-risk under the AI Act but is *not* an Annex III system, and Article 27 FRIA does not apply (Article 27(1) is limited to Article 6(2) Annex III deployers). **Segment 3 (§6)**: retail-banking deployers, European credit institutions using third-party scoring AI in their underwriting pipeline. The segment maps to Annex III §5(b), and is the only segment of the three that triggers Article 27 FRIA automatically by activity.

### 1.5 Boundary with the parent report

This annex does not re-screen any of the fifty providers in the parent report's §6. It does not produce sector aggregates that overlap the parent's §6.1–6.4. The five critical cases of the parent's §4 are referenced where relevant but not re-analysed. The unit of analysis here is the deployer segment, not the provider firm. A fourth segment (life-and-health insurance pricing deployers under Annex III §5(c)) is foreseeably the next development of the multiplier, since the parent's §4 Case E (neo-insurer with LLM chatbot and pricing pipeline) opens that thread directly; it is reserved for a v1.1 of this annex rather than included in this v1.0, which holds the "three downstream segments" scope of the title.

FIGURE 2 — THREE DEPLOYER SEGMENTS AT A GLANCE

| Dimension | Segment 1 — HRTech | Segment 2 — Healthcare | Segment 3 — Retail banking |
|---|---|---|---|
| **Deployer profile** | HR directors, CHROs of medium/large EU employers | Public hospitals, conventioned clinics, private practices | EU credit institutions (ECB-supervised or NCA-supervised) |
| **AI Act gate** | Annex III §4 (employment) | Article 6(1) via MDR/IVDR safety component | Annex III §5(b) (creditworthiness) |
| **High-risk under AIA?** | Yes (Annex III) | Yes (Art. 6(1)) | Yes (Annex III) |
| **Article 27 FRIA** | Yes for public-law deployers; conditional for private (works-council activation) | **No** — Art. 27(1) limited to Art. 6(2) Annex III | **Yes — automatically by activity** |
| **Dominant trigger** | Incident (employee / candidate / works-council letter) | Incident (Art. 50 patient question) or supervisory (MDR re-cert) | Supervisor (ACPR/EBA/ECB) + GDPR Art. 22 incidents |
| **Companion regime** | National employment law, GDPR Art. 22, works-council law | MDR/IVDR, MDCG 2025-6, GDPR Art. 9 | CRR, EBA model-risk guidelines, GDPR Art. 22, SCHUFA precedent |

# 2. Method — The path-to-discovery model

A deployer obligation that has not been examined needs an external event to be examined. This annex models that event explicitly. For each of the three segments, the analysis identifies the deployer profile, the Annex III gate triggered by the use case, the path-to-discovery (the event that surfaces the obligation), the active Article 26 paragraph subset, the Article 27 FRIA logic, the Article 25 reverse-bascule check, and a wake-up event archetype.

### 2.1 Four trigger types

Across all three segments, four trigger types account for the realistic events that surface the deployer obligation. The segments differ in which triggers dominate.

| Trigger type | Definition | Where it dominates |
|---|---|---|
| **Legal trigger** | A statutory deadline binds the deployer directly: August 2, 2026 application of Article 26, GDPR Article 35 DPIA review cycle, national workplace-AI law transposition. | All three segments at 2 Aug 2026. |
| **Supervisory trigger** | A sector supervisor sends a written request, opens an inspection, or publishes guidance that names AI controls. The deployer cannot answer the request without producing the cascade document first. | Retail banking (ACPR, ECB, EBA). Healthcare (national health authorities, Notified Bodies during MDR re-certification cycles). |
| **Contractual trigger** | A B2B counterparty (procurement, audit, due-diligence) asks the deployer in writing for its AI Act position. The deployer's own clients escalate, in turn, to their counterparties. | All three segments, with the highest velocity in HR (large-employer procurement) and banking (counterparty due-diligence). |
| **Incident trigger** | A specific person (an applicant, a patient, an employee) formally objects, complains, or initiates a procedure that requires the deployer to articulate, in writing, the basis on which the AI was used. | HR (employee complaint, works-council formal request). Healthcare (Article 50 patient transparency complaint). |

### 2.2 Segment-by-segment cascade structure

Each of the three segments below applies the same five-step structure: (1) profile of the deployer entity, (2) Annex III gate triggered, (3) path-to-discovery, which of the four trigger types dominate, (4) active Article 26 paragraph subset, (5) Article 27 FRIA logic and Article 25 reverse-bascule. The wake-up event archetype is described at the end of each segment in italics — a sentence that the deployer hears or reads, the day the obligation becomes operational rather than theoretical.

### 2.3 Article 25 reverse-bascule check

Article 25(1)(b) provides that a deployer who substantially modifies a high-risk system in a way that goes beyond the provider's intended purpose is reclassified as a provider and inherits the heavier provider obligations of Articles 16–22. The check is a check, not a finding: in most deployments, the reverse-bascule does not occur, but an organization that fine-tunes the vendor model on its own data (workforce data, patient conversations, credit files) must verify that the modification stays inside the provider's intended-purpose envelope. The three segments analysed below differ materially in how often the reverse-bascule applies.

## 3. Article 26 §1-§12 — deployer obligations reference

The table below consolidates the twelve paragraphs of Article 26 of Regulation (EU) 2024/1689 (AI Act), as published in the *Official Journal of the European Union*, OJ L, July 12, 2024. The annotations are operational summaries, not substitutes for the regulation text. The "Triggers when" column tags the deployer profile that activates each paragraph.

| § | Obligation (summary) | Triggers when |
|---|---|---|
| **§1** | Technical and organisational measures to ensure use conforms with the Article 13 provider instructions (cross-refs §3 and §6). | All high-risk deployers. |
| **§2** | Human oversight assigned to natural persons with *competence + training + authority + support*. Tested against the five capabilities of Article 14(4)(a)–(e). | All high-risk deployers. |
| **§3** | §1–§2 are *without prejudice to* other EU and national obligations (GDPR, employment law, sectoral regulation). Freedom of organisation preserved as long as oversight measures are implemented. | Transversal. |
| **§4** | Input data must be *relevant and sufficiently representative* in view of the intended purpose — only where the deployer exercises control over input data. | If deployer controls inputs. |
| **§5** | Three cumulative obligations: (a) ongoing monitoring per Article 13 notice, (b) **immediate suspension** if Article 79(1) risk detected, (c) notification of serious incident (provider → importer/distributor → market-surveillance authority). Financial-institution exception: prudential governance counts as §5 surveillance. | All high-risk deployers. |
| **§6** | Retention of automatic logs under deployer's control for *at least six months*, adapted to intended purpose, unless superseded by EU or national data-protection law. | All high-risk deployers. |
| **§7** | Prior to deployment in the workplace, employer-deployers inform workers' representatives and workers concerned. Articulation with national employment law (e.g., CSE in France, Betriebsrat in Germany). | Employer deployers · workplace use. |
| **§8** | Public authorities and EU institutions register their use in the EU database (Article 49). Deployer cannot use an unregistered high-risk system. | Public-law bodies · EU institutions. |
| **§9** | Where applicable, the Article 13 provider notice is used to meet the GDPR Article 35 DPIA obligation. Integration, not duplication. | Where DPIA already required. |
| **§10** | Post-remote biometric identification limited to law-enforcement activities. Out of scope for commercial deployers. | Law-enforcement only. |
| **§11** | Where an Annex III high-risk system takes or assists decisions concerning natural persons, those persons must be *informed* they are subject to the system's use. Cumulative with Article 50 and GDPR Articles 13–14. | Annex III · decision assistance. |
| **§12** | Cooperation with competent authorities on any measure taken toward the high-risk AI system. | Transversal. |

*Source: Regulation (EU) 2024/1689, Article 26 §1–§12, OJ L, July 12, 2024. The annotations summarise operational implications and do not substitute for the regulation text.*

# 4. Segment 1 — HRTech deployers

The parent *EU AI Act Readiness Report* §6.3 catalogues ten European HRTech *providers* (candidate evaluation, personality profiling, gamified cognitive assessments, talent intelligence with agentic workflows, CV parsing, talent sourcing, adaptive learning), with the substantial majority classified "HR + LR" — high-risk under Annex III §4 with limited-risk Article 50 transparency overlays. This annex turns to the deployers downstream of those providers. The absence of an EU-wide register of employer-deployers using HR-AI explains the typological treatment that follows: where the parent could enumerate ten provider entities, this annex profiles the segment by class rather than by name.

**METHODOLOGY NOTE**

Public registers of HR-AI deployers (employers using HRTech with AI capability) do not exist in the European Union. This segment is illustrated typologically rather than empirically. The deployer profile below is constructed from publicly disclosed HRTech-vendor client counts, large-employer procurement disclosures, and labour-law cases that have surfaced AI-assisted HR decisions. No specific employer is identified. Findings are directional.

### 4.1 Profile

The typical HR-AI deployer in this segment is a medium to large European employer, with a workforce ranging from several hundred to several tens of thousands, that has integrated one or more AI tools into the human-resources workflow. The most frequent integrations are: (a) candidate screening and ranking platforms; (b) interview-scoring or asynchronous-interview-evaluation tools; (c) employee-performance-evaluation platforms; (d) workforce-monitoring or productivity-analytics tools; (e) internal-mobility recommendation engines that assist managers in task allocation. The employer purchases the tool from an HRTech vendor; the largest global enterprise HRTech vendors publicly disclose client lists in the thousands of European enterprise customers, with aggregate Annex III §4 deployer exposure scaling accordingly. The employer integrates the tool into the existing HR information system. The employer signs a procurement contract that typically describes the vendor's compliance regime in general terms (GDPR, ISO 27001) but rarely articulates the AI Act allocation of obligations between provider and deployer.

### 4.2 Annex III §4 trigger

Annex III §4 of Regulation (EU) 2024/1689 covers AI systems used in employment, workers' management, and access to self-employment, specifically: recruitment selection, performance evaluation, task allocation, and worker monitoring. Each of the integrations listed in §4.1 above maps to one or more of these descriptors. The Article 6(3) narrow-task exception (procedural tasks, improvement of completed human activity, pattern detection without substituting for the decision, or preparatory tasks) is unlikely to apply where the system either ranks candidates or scores employees, because the system substitutes meaningfully for human decision-making at the relevant step.

### 4.3 Path-to-discovery, which triggers dominate

For HR-AI deployers, two trigger types dominate. The **incident trigger** is the most operationally consequential: a single employee, candidate, or works-council representative formally objects to an AI-assisted decision and requests, in writing, the legal basis on which the system was used. The employer must answer in writing. Standard HR documentation does not contain that answer. The cascade document is then constructed under deadline. The **contractual trigger** is the second-most frequent: a B2B counterparty (a major client, an investor in due-diligence, an audit firm) asks the employer to demonstrate AI Act compliance in its workforce processes. Legal and supervisory triggers are less frequent at present but are expected to rise after the August 2, 2026 application date and after national transpositions of works-council and labour-law provisions specific to AI-assisted decisions.

# 4. Segment 1 — HRTech deployers (continued)

### 4.4 Active Article 26 paragraph subset

Five paragraphs are typically active simultaneously for an HR-AI deployer. **§2** human oversight: the HR practitioner who reviews the AI's recommendation must satisfy the five Article 14(4) capabilities: competence to interpret the system's output, training to recognise its limits, authority to override, support to do so without operational penalty, and ability to disregard the recommendation altogether. Operational caseload caps are required: a screening team that must review fifty AI-ranked candidate files per hour cannot meaningfully exercise §2 oversight. **§5** ongoing monitoring and immediate-suspension obligation: the HR team must monitor the system for drift in selection or evaluation patterns and must suspend if an Article 79(1) risk emerges. **§6** automatic-log retention for at least six months under the deployer's control: HR-system logs already exist for GDPR purposes; the AI Act adds an inference-level granularity that may exceed what the provider supplies by default. **§7** pre-deployment information of workers' representatives and workers concerned: this is a structural deployment gate. An employer that has used the AI tool for several months without §7 information is not in compliance and cannot retroactively cure the omission. **§11** post-decision information of natural persons subject to the AI's decision: candidates not selected, employees evaluated, workers monitored — each must be informed they are subject to AI use, cumulatively with GDPR Articles 13–14.

### 4.5 Article 27 FRIA logic

Article 27(1) lists the deployer categories for which a Fundamental Rights Impact Assessment is required. The relevant categories for an HR-AI deployer are (a) "bodies governed by public law" and (b) "private entities providing public services," each deploying a high-risk Annex III system (with the exception of the §2 critical-infrastructure carve-out). The two limbs of Article 27(1) operate differently. **Limb (a)** (bodies governed by public law) applies on the basis of the deployer's legal status, without an additional test of whether the specific use case is a "public service" in the functional sense. A municipality, a regional authority, or a public agency that deploys an Annex III §4 system to manage its own personnel is therefore within the scope of limb (a) by virtue of being a public-law body, even though internal personnel management is not, in ordinary parlance, a service rendered to the public. **Limb (b)** (private entities providing public services) is read here as applying an activity-based test: the deployer must, with respect to the use case, be providing a public service. A private subcontractor operating a public-employment-service mandate on behalf of a state body falls within (b); a purely commercial private employer deploying HR-AI for its own internal recruitment falls outside (a) and (b) and does not, on activity alone, trigger Article 27. The deployer-side check is required for each entity — position-taking is not optional, and the boundary between (a), (b), and "neither" is not always self-evident.

### 4.6 Article 25 reverse-bascule check

Most HR-AI deployers do not fine-tune the vendor model on workforce data. The typical deployment is configuration of weights, customisation of decision thresholds, and integration with the existing HR information system, not retraining. Where the vendor offers a fine-tuning option and the employer accepts it, the reverse-bascule under Article 25(1)(b) becomes a serious risk. The principal scenario is the employer who fine-tunes the vendor's screening model on its own historical hiring data: the modification is likely to change the system's intended purpose in ways the vendor's technical file does not anticipate, and the employer may inherit Article 16 provider obligations.

> **Wake-up event archetype.** A works-council representative writes a formal request to the HR director: "We were not informed before this tool was deployed. Please provide the documentation describing how the system reaches its recommendations, the human oversight in place, and the legal basis under the AI Act and applicable national employment law." The HR director's standard reply ("the vendor is GDPR-compliant") does not answer the request. The cascade document (Article 26 §2, §5, §6, §7, §11 mapped, FRIA scope memo, §7 information record) is what answers the request.

# 5. Segment 2 — Healthcare deployers

### 5.1 Profile

The typical healthcare deployer in this segment is a European hospital, clinic, or medical practice that has integrated an AI-driven clinical-documentation tool. Most often, a real-time scribe application that listens to patient-clinician conversations and produces structured clinical notes (into its everyday workflow. Three sub-types are common across the European healthcare landscape: (a) public hospitals operating under national health systems (CHU in France, Krankenhaus public in Germany, NHS trusts in the United Kingdom for activities still in scope of EU compliance, public hospitals in the Netherlands and across the Nordic systems); (b) conventioned private clinics) private institutions that deliver care reimbursed by national or social-security systems and that, in many member states, qualify as providers of a public service; (c) purely commercial private practices, where the deployer is a private medical professional or a private clinic operating outside the public-service framework. The deployer purchases or licenses the AI scribe from a medical-AI provider, integrates it into the consultation workflow, and treats compliance as covered by the provider's clinical CE marking and the existing GDPR Article 9 framework.

### 5.2 Article 6(1) trigger via MDR/IVDR safety-component pathway

Article 6(1) AIA imposes two cumulative conditions: (a) the AI system is a safety component of a product covered by Union harmonisation legislation listed in Annex I (MDR or IVDR), or is itself such a product; *and* (b) that product is required to undergo a third-party conformity assessment in view of its placing on the market. For a clinical-scribe application qualified as Medical Device Software (MDSW), the second condition depends on the MDR/IVDR risk class. **MDR Class IIa, IIb, or III** requires Notified Body assessment and therefore satisfies condition (b); the AI Act high-risk classification under Article 6(1) follows. **MDR Class I non-sterile, non-measuring, non-reusable surgical** is self-certified and does not engage Notified Body involvement; condition (b) is not met, and Article 6(1) does not trigger. **MDR Class I sterile, measuring, or reusable surgical** requires Notified Body involvement for the relevant aspect (sterilisation, metrology, or reprocessing); condition (b) is satisfied and Article 6(1) triggers. The mapping is consolidated in MDCG 2025-6 / AIB 2025-1 (June 19, 2025), Table 1. The Article 6(1) trigger therefore applies to the substantial majority, but not the totality — of MDSW clinical scribes on the market; the residual self-certified Class I scribes follow a different path (typically Article 50 transparency only).

Annex III §5(a) addresses AI systems used to evaluate the eligibility of natural persons for essential public-assistance benefits and services, which is a different category from clinical documentation and does not apply to the MDSW scribe pathway. Where the underlying technology is a frontier general-purpose AI model embedded as the language backbone, the deployer's interaction with end users is also subject to Article 50; the GPAI provider's Article 53 documentation reaches the deployer indirectly via the Article 13 instructions for use.

*Two adjacent Annex III categories may apply to other healthcare deployers but are out of scope of this segment: §5(a) (AI used by public authorities to evaluate benefit eligibility, and §5(d)) AI used in emergency dispatch and patient triage. Both are Annex III high-risk and trigger Article 27 FRIA conditionally; this annex addresses the more numerically significant Article 6(1) MDSW pathway, in its NB-assessed configuration.*

### 5.3 Path-to-discovery, which triggers dominate

For healthcare deployers, three trigger types dominate. The **incident trigger** is most often a patient who, having understood that an AI listened to their consultation, formally requests the legal basis under which the AI was used and the documentation of how the AI's output entered their medical record. The Article 50 transparency obligation makes this request reasonable. The deployer must answer in writing. Standard hospital documentation does not contain the answer. The **supervisory trigger** arrives during the next MDR/IVDR re-certification cycle, where the Notified Body or national health authority asks the hospital to evidence the AI Act overlay onto its existing medical-device quality management system. The **legal trigger** is the August 2, 2026 application date itself: hospitals operating as public-law bodies face Article 27 FRIA notification under §3 and Article 49 EU-database registration under §8 from that date.

# 5. Segment 2 — Healthcare deployers (continued)

### 5.4 Active Article 26 paragraph subset

Five paragraphs are typically active simultaneously for a healthcare deployer using an AI scribe qualified as MDSW in clinical workflow. **§1** conformant use: the deployer must implement technical and organisational measures to ensure the system is used in line with the provider's Article 13 instructions for use and the MDR/IVDR clinical-evaluation envelope. **§2** human oversight: the clinician retains diagnostic judgment and is the natural-person operator under Article 14(4); the AI scribe is not a clinical decision-maker and the §2 oversight architecture must reflect that. The clinician's burden is not "approve the scribe's output" but "verify before signing." **§5** ongoing monitoring and immediate-suspension if a §79(1) risk is detected; the MDR/IVDR post-market surveillance regime provides part of the §5 evidence base but does not absorb it. **§6** log retention under deployer's control for at least six months, distinct from the MDR Eudamed vigilance logs that the medical-device manufacturer maintains; the deployer-side log captures the inference-level events specific to each consultation. **§9** integration of the Article 13 provider notice into the existing GDPR Article 35 DPIA, recognising that GDPR Article 9 high-risk processing already exists for clinical data and that the AI Act overlay does not duplicate but complements. **§12** cooperation with competent authorities, particularly during MDR/IVDR re-certification cycles where Notified Body inquiries are foreseeable.

Two paragraphs that are *not* activated for the Article 6(1) clinical-scribe pathway are worth naming explicitly, because they are commonly assumed to apply by analogy. **§11** (information of natural persons subject to a high-risk AI system) is by its terms restricted to Annex III deployers; Article 6(1) systems are not Annex III. Patient information about the AI's involvement in the consultation is therefore based on Article 50 transparency, GDPR Articles 13–14, and applicable national patient-rights law, rather than on Article 26 §11. The substantive obligation to inform the patient is not weaker; the legal basis is different. **§7** (workplace deployment information of workers' representatives) does not apply because the patient is not the deployer's worker.

### 5.5 Article 27 FRIA logic

Article 27(1) limits the FRIA obligation to deployers of "high-risk AI systems referred to in Article 6(2)" (i.e. Annex III systems. Article 6(1) safety-component systems, including the MDSW clinical scribe in scope of this segment, are *not* covered by Article 27. A public hospital deploying an AI scribe under the Article 6(1) pathway is therefore not, on activity alone, required to perform a FRIA under the AI Act, even though it qualifies as a public-law body. This is a deliberate boundary set by the regulation's text and should not be assumed away. Where the same hospital separately deploys an Annex III system, for example, a §5(a) eligibility-evaluation system to determine access to specialised care, or a §5(d) emergency-triage system — Article 27 FRIA applies to that separate deployment by activity for §5(d) automatic triggers, or by public-law-body status for §5(a) or other Annex III categories. The Article 27 FRIA boundary is therefore drawn around the specific use case, not around the deployer entity.

# 5. Segment 2 — Healthcare deployers (continued)

### 5.6 Article 25 reverse-bascule check

Healthcare deployers occasionally fine-tune the vendor's AI scribe on their own corpus of patient conversations to improve specialty-specific accuracy, terminology recognition, or local-language adaptation. Where this occurs, the modification is likely to constitute a substantial modification under Article 3(23) and triggers Article 25(1)(b) reverse-bascule into provider status. The hospital that fine-tunes does not, in most cases, intend to become a provider. The technical-documentation, post-market-monitoring, and quality-management-system obligations that follow are not absorbed by the vendor. The reverse-bascule check is therefore operationally consequential at the moment a hospital negotiates a fine-tuning option with its scribe vendor. In the MDR/IVDR context specifically, a substantial modification under the AI Act may also constitute a "significant change" under MDR Article 120(3) or its IVDR equivalent, with separate notified-body implications; the AI Act and MDR/IVDR substantial-modification tests are independent and must be assessed in parallel.

FIGURE 3 — ARTICLE 25(1)(B) REVERSE-BASCULE (HEALTHCARE WORKED EXAMPLE)

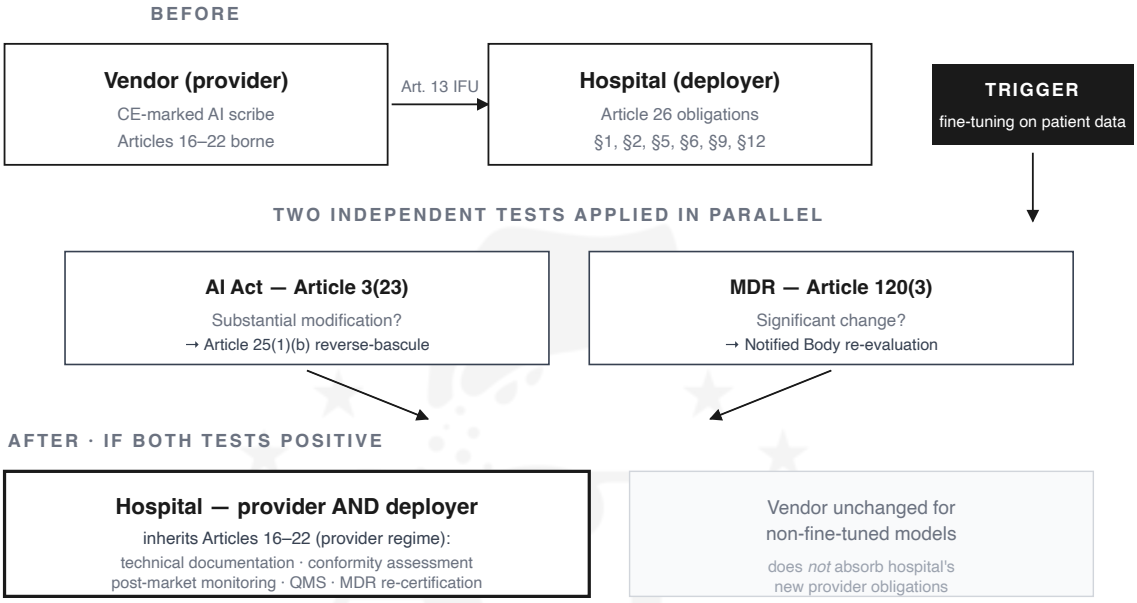


**Reading.** The two tests (AI Act Art. 3(23); MDR Art. 120(3)) are independent and must be applied in parallel. A hospital that contractually accepts fine-tuning on its own patient corpus enters the reverse-bascule path at that moment, not retroactively. The check belongs in the procurement workflow, not in post-deployment audit. Articles 16–22 cannot be delegated back to the vendor by contract.

> **Wake-up event archetype.** A patient writes: "I noticed an AI listened to our consultation. Please tell me on what legal basis the AI was used, what was recorded, how long it is retained, who has access, and whether its notes informed clinical decisions about me." Article 50 transparency makes this reasonable. The cascade document (Article 26 §1, §2, §5, §6, §9, §12 (Article 6(1) pathway), the Article 50 + GDPR 13–14 patient record, and the Article 25 boundary log) answers the letter.

# 6. Segment 3 — Retail-banking deployers

### 6.1 Profile

The typical retail-banking deployer in this segment is a European credit institution licensed under the Banking Union prudential framework (supervised either directly by the European Central Bank for significant institutions or by the national competent authority (in France, the Autorité de Contrôle Prudentiel et de Résolution; in Germany, BaFin; in the Netherlands, De Nederlandsche Bank; in Italy, Banca d'Italia; in Spain, the Banco de España)) that has integrated a third-party AI scoring model into its underwriting pipeline. The deployer is the credit institution, not the scoring vendor. The typical configuration is: the institution licenses a vendor scoring engine, configures decision thresholds, integrates the scoring output into the loan-officer interface, and retains the underwriting decision in the loan officer's hands subject to operational caseload pressures. Under the Capital Requirements Regulation and the EBA's existing ICAAP and model-risk-management frameworks, the institution already operates a model-risk-governance regime; the AI Act layer is added on top and is not absorbed by it.

### 6.2 Annex III §5(b) trigger

Annex III §5(b) of Regulation (EU) 2024/1689 covers AI systems intended to evaluate the creditworthiness of natural persons or to establish their credit score. Retail consumer-credit underwriting in a licensed credit institution falls squarely within this category. The Article 6(3) narrow-task exception is not available where the system substantively contributes to the credit decision, which is the typical configuration. The §5(b) gate is therefore activated for the great majority of retail-credit deployers using vendor AI scoring. Where the underlying technology integrates a frontier general-purpose AI model, for example, for structured-document summarisation feeding the underwriting file — the GPAI provider's Article 53 documentation reaches the institution indirectly via the Article 13 instructions for use, and Article 50 transparency obligations apply directly to any deployer-side interaction with applicants.

### 6.3 Path-to-discovery, which triggers dominate

For retail-banking deployers, three trigger types dominate, and the timing is unusually compressed. The **supervisory trigger** is structural: ACPR, ECB through the Joint Supervisory Team for significant institutions, or the national competent authority routinely inspects model-risk-management arrangements. The supervisor's first AI-Act-aware inspection cycle is foreseeable in the months following the August 2, 2026 application date, and the institution that has not articulated the deployer-side cascade in advance answers the inspection in real time. The **legal trigger** is automatic: Article 27(1) names §5(b) deployers explicitly, which means the FRIA is automatic by activity for every retail-credit institution using AI scoring — independent of the public-law-body test. The **incident trigger** arrives in the form of an applicant whose loan was refused or whose terms were adverse, who exercises GDPR Article 22 rights against automated decision-making and who, upon receiving the institution's response, escalates to the national data-protection authority, the national consumer ombudsman, or (in a growing number of jurisdictions) initiates litigation (see CJEU Case C-634/21, *SCHUFA Holding AG*, judgment of December 7, 2023, on the scope of automated decision-making under GDPR Article 22 in the credit-scoring context).

# 6. Segment 3 — Retail-banking deployers (continued)

### 6.4 Active Article 26 paragraph subset

Five paragraphs are typically active simultaneously for a retail-banking deployer using third-party AI scoring. **§2** human oversight: the loan officer is the natural-person operator and must satisfy the five Article 14(4) capabilities. Operational caseload caps are critical: an underwriting team that processes hundreds of files per day cannot exercise meaningful §2 review unless caseload is calibrated to permit substantive examination of each AI-ranked recommendation. **§5** ongoing monitoring and immediate-suspension upon §79(1) risk; the recital-noted prudential-governance exception does not eliminate the §5 obligation but allows the institution's existing model-risk-monitoring regime to count toward §5 surveillance where it covers the relevant AI-specific risks. **§6** automatic-log retention for at least six months under the institution's control, distinct from the model-validation logs already kept under the model-risk-management regime; the deployer-side log captures inference-level events including each scoring decision and its associated input. **§9** integration of the Article 13 provider notice into the GDPR Article 35 DPIA already required for credit-scoring processing, and articulation with GDPR Article 22 on solely automated decision-making with significant effect — the institution must determine, for each scoring use case, whether the scoring is "solely automated" within the meaning of Article 22 and what the appropriate human-intervention safeguard is. **§11** applicant information: applicants whose creditworthiness is evaluated must be informed they are subject to AI use, cumulatively with the GDPR Article 13–14 information already provided at the point of application.

### 6.5 Article 27 FRIA logic

The FRIA is automatic for retail-banking deployers of AI credit scoring because Article 27(1) names §5(b) deployers explicitly. The institution is not entitled to wait for a supervisory request; the FRIA is required to be in place before deployment or before the next material change of the system. Article 27 §3 mandates pre-deployment notification to the national market-surveillance authority for the AI Act. Article 27(4) is explicit that FRIA complements the GDPR DPIA and does not replace it. The institution must therefore produce three documents that are operationally distinct and substantively interoperable: the existing model-risk-management documentation, the GDPR Article 35 DPIA, and the Article 27 FRIA.

### 6.6 Article 25 reverse-bascule check

Retail-banking deployers rarely fine-tune the vendor scoring model on proprietary credit files in the conventional sense. The typical institution configures decision thresholds, calibrates the score to its risk-appetite framework, and integrates the scoring into the underwriting workflow without altering the model weights. Where the institution does fine-tune, for example, by retraining the scoring model on its own historical book to improve calibration to its specific clientele — the modification is likely to constitute a substantial modification under Article 3(23) and triggers Article 25(1)(b) reverse-bascule into provider status. The institution that fine-tunes inherits Article 16 provider obligations and is responsible for the Article 9 risk-management system, the Article 11 technical documentation, and the post-market monitoring of its now-proprietary scoring model. The reverse-bascule is therefore the most consequential boundary check in this segment.

> **Wake-up event archetype.** A letter from the national competent authority or the Joint Supervisory Team: "We are conducting a horizontal review of AI-supported credit decisions. Please provide, for each scoring use case, the Article 26 paragraph map, the Article 27 FRIA scope, the Article 25 substantial-modification position, and the integration with your ICAAP framework. Response expected within thirty business days." The institution that has not previously produced the cascade discovers that it is not optional.

# 7. Cross-segment patterns

The three segments analysed above differ materially in profile, in the dominant trigger type, and in the Article 27 FRIA logic. They converge on a small number of shared structural features that the deployer should recognise irrespective of the segment.

### 7.1 The shared base cascade

Across all three segments, the active Article 26 paragraph subset converges on §2 (human oversight), §5 (ongoing monitoring), and §6 (log retention) as the operational base cascade for any deployer of a high-risk system. §11 (information of natural persons subject to the AI Act high-risk system) is by its terms restricted to deployers of Annex III systems and is therefore active for HR (§4) and banking (§5(b)) but not for the Article 6(1) MDR/IVDR clinical-scribe pathway analysed for healthcare; in the Article 6(1) case, the equivalent information obligation flows from Article 50 transparency, GDPR Articles 13–14, and applicable national patient-rights law. The remaining paragraphs add layers — §7 information of workers becomes structural in employment, §8 EU-database registration applies to public-law deployers of Annex III systems, §9 DPIA integration applies wherever GDPR Article 35 obligations already exist, §12 supervisor cooperation becomes pressing where a sector supervisor exists.

### 7.2 The dominant trigger map

The trigger map across the three segments is asymmetric. HR deployers are most often triggered by the incident type — an employee, a candidate, a works-council representative initiating a written request. Healthcare deployers are most often triggered by either an incident (Article 50 patient question) or a supervisory event (next MDR re-certification cycle). Retail-banking deployers face the most compressed timeline: the supervisory trigger is structural and predictable, the legal trigger is automatic by activity (Article 27(1) §5(b)), and the incident trigger via GDPR Article 22 is increasing in volume. The implication is operational: the segment with the most predictable trigger map (banking) is also the segment that has the least time to respond to a trigger when it arrives.

### 7.3 The provider-versus-deployer self-classification confusion

An organization that purchases an AI tool and configures it for use under its own authority is a deployer within the meaning of Article 3(4), regardless of how the organization labels itself in its own documentation. The most frequent confusion observed across the three segments is the deployer that describes itself as "the user" or "the customer" of the AI system and concludes that compliance is therefore the provider's problem. The Article 26 obligations are not derived from a label; they are derived from the function the organization performs vis-à-vis the AI system. The deployer that configures, monitors, and uses the system is the deployer whether it calls itself so or not.

### 7.4 The Article 25 reverse-bascule rate is low but consequential

Across all three segments, the rate of reverse-bascule under Article 25(1)(b) is low — most deployers do not fine-tune the vendor model on their own data. Where it occurs, the consequences are consequential: the deployer inherits the provider regime (Articles 16–22), including the technical-documentation, conformity-assessment, and post-market-monitoring obligations that the deployer is rarely organisationally equipped to absorb. The boundary check should be performed at the moment a fine-tuning option is contractually accepted, not retroactively.

## 7. Cross-segment patterns (continued)

### 7.5 Silence does not propagate symmetrically

The parent report's §3.1 observes that the provider's compliance deficit is inherited by the deployer. The reciprocal does not hold symmetrically: a deployer that has produced its own cascade does not regularise the provider's underlying defects. Where the provider's technical file is missing, incomplete, or non-conformant, the deployer's well-prepared cascade still rests on a defective foundation. The deployer's path forward in that situation is documentary — request the missing items from the provider in writing, escalate where they are not forthcoming, and document the gap if it persists. Silence at the provider level does not release the deployer from §1 (technical and organisational measures to ensure conformant use); it raises the bar on it.

### 7.6 The deployer's seat is rarely held by one function

The deployer obligation under Article 26 is borne by the legal person, but the operational responsibility is distributed across functions: legal, compliance, IT, business operations, and (depending on the segment) works-council representation, clinical leadership, or model-risk-management. None of these functions, on its own, has the authority or the visibility to produce the cascade document. An organization that recognises itself in one of the three segments above and that wants to be ready by August 2, 2026 typically appoints a single accountable owner for the cascade, with cross-functional reporting. The cascade document is produced under that owner's authority, not as an artefact of any single function.

### 7.7 The cascade is dated by design

The cascade document carries its date with it. A document produced on August 1, 2026 and a document produced on September 2, 2026 in response to a supervisor inquiry both meet the substantive Article 26 obligation when the content is equivalent. *Whether* the supervisory weight attached to the two dates differs in practice (proactive versus reactive posture) is a reasonable inference rather than an established fact: no decided case, supervisory guidance, or AI Office communication has yet articulated a formal dating asymmetry, and supervisors are bound by their statutory mandates regardless of timing. The annex's recommendation to begin work earlier rather than later rests on the practical observation that documents produced under a thirty-day inquiry window tend to be less complete than documents produced under self-directed timing, not on a juridical claim about supervisory deference.

# 8. From multiplier to action

The parent report's §8 observes that "Classification comes first. Everything else follows from it." This annex develops what "everything else" looks like for the three deployer segments. The deployer that recognises itself in one of the segments above and that wants to be ready by August 2, 2026 produces, between recognition and the deadline, four operational outputs.

### 8.1 Output one — the Article 26 paragraph map

For the specific use case in scope, the deployer maps which of the twelve Article 26 paragraphs are active, which are inactive, and the operational artefact that evidences each active paragraph. The output is a one-to-three-page document, dated, that lists the paragraphs and points to the underlying operational evidence (oversight architecture, monitoring procedure, log-retention configuration, information notice, DPIA integration). The map is not the evidence; it is the index of the evidence.

### 8.2 Output two — the FRIA scope memo

For deployers triggering Article 27 (automatic by activity for retail-banking §5(b) within the segment scope of this annex, conditional by public-law-body or public-service status for HR (limbs (a) and (b) of Article 27(1))) the FRIA scope memo articulates which Article 27(1) limb applies, what the FRIA covers and does not cover, the relationship to the existing GDPR Article 35 DPIA, and the Article 27 §3 pre-deployment notification status. For deployers *not* triggering Article 27, including all Article 6(1) safety-component deployments such as the clinical-scribe MDSW pathway analysed in §5 — the memo records the conclusion (with reasoning) that no FRIA is required, which is itself a position-taking the deployer should be able to defend in writing.

### 8.3 Output three — the Article 25 boundary log

The boundary log records, at a defined point in time, the conclusion that the deployer's configuration choices (decision thresholds, integration parameters, fine-tuning if any) do or do not constitute a substantial modification under Article 3(23). The log is updated each time a material configuration change is made. Where the log concludes that the boundary is approached but not crossed, the reasoning is preserved so that, in the event of supervisory inquiry, the deployer can demonstrate the boundary check was performed in advance rather than reconstructed retrospectively.

### 8.4 Output four — Article 49 EU-database registration where applicable

Public-law-body deployers and EU institutions register their use in the EU database under Article 49, by reference to Article 26 §8. Registration is not a self-service portal in May 2026; the deployer should plan the registration on the assumption that the database opens for deployer entries on or shortly before August 2, 2026 and that early-week traffic will be high. Pre-2-August preparation of the registration content (system identification, intended-purpose statement, deployer entity identification) saves time on the day.

### 8.5 What the four outputs together produce

The four outputs together constitute the dated, article-mapped document that the parent report's §8 identifies as the missing artefact. They are not a single document; they are four documents that interoperate. A deployer that has produced them is not "compliant" — compliance is a regulator's determination on a specific record. But the deployer that has produced them has answered the four questions a supervisor, a counterparty, an employee, a patient, or an applicant is reasonably likely to ask in the months following August 2, 2026.

## 9. Positioning

Sprinkling Act operates as a pre-conformity advisory firm for European deployers. The work focuses on the deployer-side cascade: use-case classification, Article 26 paragraph mapping, Article 27 FRIA scoping where applicable, and Article 25 boundary documentation. The firm does not develop AI systems, does not sell governance software, and does not certify compliance.

**Verifiable references.** Parent *EU AI Act Readiness Report* (April 2026) (DOI **10.5281/zenodo.19671329**, indexed on Zenodo (CERN). Author identifier) ORCID **0009-0002-5093-8550**. Founder profile (**sprinklingact.com/about/lamar-shucrani**. Sectoral pages) **sprinklingact.com/healthtech** (MDR × AIA interplay for medical AI). All references are independently verifiable; no commercial relationship is required to retrieve them.

**What this annex offers the deployer who recognises themselves in it.** A free initial diagnostic that surfaces which of the six regulatory gates apply to a specific use case, what the Article 26 paragraph subset is, and whether Article 27 is automatic, conditional, or inapplicable, is available at **sprinklingact.com/check**. A personalised cascade document (Article 26 paragraph map, FRIA scope memo, Article 25 boundary log, and where applicable Article 49 registration content) is the firm's principal commercial output and is delivered as a dated, hash-stamped report. The diagnostic is the entry point; the report is the operational artefact.

## 10. References

**Primary regulatory source.** Regulation (EU) 2024/1689 of the European Parliament and of the Council of June 13, 2024, laying down harmonised rules on artificial intelligence (the "AI Act"), *Official Journal of the European Union*, OJ L, July 12, 2024. Particular reference: Articles 3(4), 3(23), 3(63), 5, 6, 13, 14(4), 25, 26 §1–§12, 27 §1–§5, 49, 50, 51, 53, 79(1).

**Adjacent regulatory sources.** Regulation (EU) 2017/745 (Medical Devices Regulation (MDR). Regulation (EU) 2017/746) In Vitro Diagnostic Medical Devices Regulation (IVDR). Regulation (EU) 2016/679 — General Data Protection Regulation (GDPR), particular reference to Articles 13–14, 22, 35.

**Sectoral guidance.** MDCG 2025-6 / AIB 2025-1, "Interplay between MDR & IVDR and AIA," Medical Device Coordination Group + Artificial Intelligence Board, June 19, 2025 (Table 1 (MDR/IVDR classification mapping to Article 6(1) AIA high-risk). MDCG 2019-11 rev.1, "Qualification and classification of software) Regulation (EU) 2017/745 and 2017/746," MDCG, revised June 2025. EBA guidelines on internal governance and on the management of model risk for credit institutions, in their successive consolidated versions applicable in 2026.

**Secondary reference.** Kenney, Noah M., *Governing Intelligence: Law, Privacy, Security, and Compliance in the Age of Artificial Intelligence*, Digital 520, First Edition 2026 — referenced as a publicly available example of an article-mapped overview of EU AI Act obligations across the provider-deployer boundary.

Correspondence: **contact@sprinklingact.com**.

## 11. Limitations

This annex is an analytical reference, not a conformity assessment. Specific classifications for any specific deployer require a tailored screening. The three segment profiles are constructed for illustrative purposes; the HRTech segment in particular is illustrated typologically, owing to the absence of public registers of HR-AI deployers in the European Union. No specific employer, hospital, clinic, credit institution, or other organization is identified by name in this annex; the descriptors used (sub-types, sizes, configurations) are drawn from publicly disclosed information about the relevant sectors and from the parent report's §6 sector breakdown.

The "more than 3,500 European organizations" figure adopted from the parent report's §3.1 is an estimate derived from publicly stated client counts of four named high-risk providers in the parent report's sample; the underlying disclosures are subject to the providers' own reporting choices and may not be fully comparable across providers. The "tens of thousands" and "hundreds of thousands" extrapolations are orders of magnitude rather than direct measurements.

The Digital Omnibus negotiation status reported in §0.8 reflects the position of the European Parliament and the Council as of May 2026; that position may evolve before final adoption.

The classification statements reflect the author's reading of Regulation (EU) 2024/1689 (OJ L, July 12, 2024) and of MDCG 2025-6 / AIB 2025-1 (June 19, 2025). In particular: Article 27 FRIA is limited to Article 6(2) Annex III deployers and does not apply to Article 6(1) safety-component systems; Article 26 §11 is restricted to Annex III deployers. These textual limits drive the Article 6(1) versus Annex III distinction in segment 2. Final determinations for any specific deployer require qualified legal counsel.

Note on parent-report consistency: parent §9.4 Sample Composition lists "FR, CH, IL, EU" for the FinTech / InsurTech sub-sample, where "IL" stands for Israel. The parent's cover and §1 Executive Summary describe the report as covering 50 European AI companies across 15+ countries (FR, DE, BE, NL, PL, DK, CH, IE, UK), and Israel is not in that list. The FinTech entry 19 in parent §6.2 carries an "IL/FR" descriptor reflecting an Israeli acquisition origin with French deployment. A future revision of the parent will clarify the inclusion of an Israeli-origin entity within the European-sample frame; this annex inherits the parent's sample composition as published and signals the inconsistency for transparency rather than for re-derivation.

This annex does not constitute legal advice, regulatory advice, or any other professional advice. Sprinkling Act is not a law firm, not a Notified Body, not a certification body, and is not affiliated with the European Commission, the European Parliament, the AI Office, the Artificial Intelligence Board, or any national supervisory authority. Organizations recognising themselves in one of the three segments should consult qualified legal counsel and, where applicable, the relevant sector supervisor or Notified Body before making compliance decisions.

## 12. Document metadata

| Field | Value |
|---|---|
| **Annex version** | 1.0 — first edition |
| **Drafting period** | April 15 – May 6, 2026 (window opened three days after parent publication) |
| **Publication date** | May 6, 2026 |
| **Parent document** | EU AI Act Readiness Report (Sprinkling Act, April 12, 2026) |
| **Parent DOI** | 10.5281/zenodo.19671329 |
| **Annex DOI** | 10.5281/zenodo.20042175 |
| **License** | Creative Commons Attribution 4.0 International (CC-BY 4.0) — free to download, share, and cite with attribution |
| **Suggested citation** | Shucrani, L. B. (2026). *EU AI Act Readiness — Annex A: The Deployer Multiplier (May 2026)*. Zenodo. https://doi.org/10.5281/zenodo.20042175 |
| **Author ORCID** | 0009-0002-5093-8550 |
| **Badge code** | SA-20260503-0001 · verifiable at sprinklingact.com/verify/SA-20260503-0001 |
| **Analyst** | Lamar B. Shucrani, Founder, Sprinkling Act. Brussels-based AI Act analyst working on the pre-conformity layer for European deployers; author of the parent *EU AI Act Readiness Report* (April 2026, Zenodo DOI 10.5281/zenodo.19671329). |
| **Cryptographic timestamp** | OpenTimestamps detached proof (.ots) anchored to Bitcoin via four calendar servers (a/b OpenTimestamps pools, Eternity Wall, Catallaxy). Mirror: github.com/sprinkling-act/timestamps |
| **Errata and corrections** | Reported transparently via versioned re-publication (v1.1, v1.2, …) with change log; the .ots proof of each edition remains valid for that edition. |
| **Independence** | No commercial relationship existed, at the time of publication, between Sprinkling Act and any organization in the three segments analysed in this annex, beyond what is disclosed in the parent report's §3 in respect of the parent's own screened sample. |

## Most deployers do not know they are deployers.

The cascade from a small number of high-risk providers to thousands of European deployer organizations is the AI Act's most actionable structural finding. This annex names the three downstream segments where that cascade is most visible, and what each deployer produces between recognition and August 2, 2026.

**sprinklingact.com/check**

9 questions. 60 seconds. Zero data collected. Free.