REFERENCE CARD · CC BY 4.0

Sprinkling Act Methodology Card

April 2026 edition (v1.1)

CONTENTS

Card details	
Edition	April 2026 (v1.1)
Regulatory freeze	March 2026
Publication date	April 27, 2026
Source of authority	Regulation (EU) 2024/1689 — OJ L, July 12, 2024
Companion documents	Parent screening report (Zenodo DOI 10.5281/zenodo.19671329); Annex A (Zenodo DOI 10.5281/zenodo.20042175)
License	Creative Commons Attribution 4.0 International (CC BY 4.0)
Author	Lamar B. Shucrani · ORCID 0009-0002-5093-8550
Independence	Not a law firm, not a Notified Body, not a certification body

0. Purpose and scope

0.1 What this card is

This card is the methodological reference for Sprinkling Act position assessments under Regulation (EU) 2024/1689 (the EU AI Act). It documents the six-gate decision logic, the article-mapped output format, the verification protocol, and the alignment with international risk-management standards. It is the reference cited in Sprinkling Act published research (the *EU AI Act Readiness Report* of April 2026 and *Annex A — The Deployer Multiplier* of May 2026).

0.2 What this card is not

This card is not legal advice, not a conformity assessment within the meaning of Article 43 AIA, and not a certification of any kind. It is not a substitute for the qualified legal counsel that an organisation should consult for binding determinations, and it is not a substitute for the Notified Body assessment that Articles 43 and 44 require for certain categories of high-risk AI systems. The card describes a framework that an organisation can adopt, contest, or extend; it does not describe the only defensible reading of the AI Act.

0.3 Scope of application

The framework applies to AI systems and general-purpose AI models within the scope of Regulation (EU) 2024/1689 as defined in Article 2, deployed or made available in the European Union, regardless of the deployer's or provider's place of establishment. The framework does not cover national implementing legislation (which varies by Member State), sector-specific regulations beyond those expressly cross-referenced in the AI Act (notably MDR, IVDR, and Annex I product legislation), or the interaction with the GDPR beyond the points where the AI Act explicitly cross-references it (Articles 10(5), 26 §9, 50, 79).

0.4 Position in the Sprinkling Act corpus

This card sits at the methodological centre of three publicly accessible artefacts. The **parent *EU AI Act Readiness Report*** (April 12, 2026, Zenodo DOI 10.5281/zenodo.19671329) applies this card to fifty European AI providers and produces sector-level findings; §3.1 of the parent identifies the deployer-cascade phenomenon. **Annex A — The Deployer Multiplier** (May 6, 2026, Zenodo DOI 10.5281/zenodo.20042175) develops that cascade for three downstream segments (HRTech, Healthcare, Retail-banking). This card is the underlying gate logic that both apply; it is the citable reference that allows a third-party reviewer to inspect the framework independently of any single screened firm.

0.5 What you will find in this document

The card is organised around twelve substantive sections plus closing metadata. §1 names the six core principles. §2 places this card in the Sprinkling Act published corpus. §3 details the six regulatory gates one by one. §4 maps the assessment to the article-level obligations of Articles 9, 10, 14, and 15 (the four obligations most relevant to the bias-and-fairness frontier). §5 decomposes Annex III into its eight categories with their classification consequences. §6 covers the GPAI obligations regime (Articles 50, 51, 53, 55). §7 names the deployer-side obligations of Articles 26 and 27 and points the reader to *Annex A* for the developed cascade. §8 describes the output format of an assessment. §9 documents the verification protocol (OpenTimestamps, badge code, audit trail). §10 maps the framework onto international risk-management standards. §11 names the limitations. §12 lists the primary references.

1. Six core principles

The framework rests on six methodological commitments that are visible in the output of every assessment. They are the operational test of whether a given assessment was produced by this framework or by an adjacent one.

1.1 Article mapping

Every classification cites a specific article, paragraph, or annex of Regulation (EU) 2024/1689 in the published version of the Official Journal (OJ L of July 12, 2024). The framework does not interpret beyond the regulatory text, beyond the published Commission guidance (notably MDCG 2025-6 / AIB 2025-1 of June 19, 2025 on the MDR-IVDR-AIA interplay), and beyond the published Court of Justice decisions that bind on the questions in scope (notably SCHUFA Holding C-634/21 of 7 December 2023 for credit-scoring under GDPR Article 22). Where interpretation is required, the framework names the interpretation as such and recommends qualified legal counsel.

1.2 Gate logic

The six gates are evaluated in sequence. The first gate triggered determines the final classification. A trigger at Gate 01 (Article 5 prohibited practices) ends the assessment: the system cannot be deployed legally, and no further gate is evaluated. Otherwise, the gates are evaluated in order from Gate 02 (Article 6(1) safety component) through Gate 06 (Article 53 GPAI standard obligations); each gate's outcome is documented even when the gate is not triggered, so that a third-party reviewer can reproduce the path.

1.3 Audit trail

Every classification carries a traceable path. The path is gate-by-gate and article-by-article. Each step is documented in the assessment artefact such that a third-party reviewer can reproduce the classification using only the published methodology card, the published primary sources, and the inputs the deployer or provider has supplied. The artefact is the unit of verification, not the conclusion.

1.4 Temporal stability indicator

Every classification carries a stability indicator: **STABLE**, **MODERATE**, or **UNSTABLE**. STABLE indicates that the classification rests on settled regulatory text, settled Commission guidance, and binding case law; the classification is not expected to move under foreseeable evolutions. MODERATE indicates that the classification rests on regulatory text but on guidance that may evolve (notably the AI Office's progressive output on Annex III categories and on Article 6(3) exception narrowness). UNSTABLE indicates that the classification rests on a frontier where guidance has not yet been issued, or where the Court of Justice has not yet ruled, and where the classification could move materially as guidance is issued. The Digital Omnibus on AI, currently in trilogue, is treated as MODERATE volatility for the deadlines it would shift; the underlying classifications are not affected.

1.5 No interpretation in the client's favour

Where classification is uncertain, the uncertainty is documented in the artefact and explicitly flagged with a recommendation to seek qualified legal counsel. The framework does not select the lower-risk reading where two readings are equally defensible; it names both readings and lets the client, the client's counsel, or the supervising authority resolve the choice. This commitment is the operational counterpart of the framework's independence: there is no commercial advantage to under-classifying the system.

1.6 Dual versioning

Every artefact carries the methodology edition (this card, e.g. April 2026 v1.1) and the regulatory freeze date (the date through which regulatory developments have been incorporated). An assessment produced on April 25, 2026 against the April 2026 / regulatory-freeze-March-2026 envelope is a different artefact from one produced on June 25, 2026 against a June 2026 / regulatory-freeze-May-2026 envelope. Both are reproducible against their own envelope; neither is "right" or "wrong" outside its temporal frame.

2. Position in the Sprinkling Act published corpus

This card is one of three publicly accessible methodological artefacts. The three are intended to be read together by a reviewer who wants to verify the framework: this card establishes the gate logic; the parent screening report applies it to a sample; the annex develops the cascade onto the deployer side.

2.1 Parent EU AI Act Readiness Report (April 2026)

The parent report (Sprinkling Act, April 12, 2026; Zenodo DOI 10.5281/zenodo.19671329, license CC BY 4.0) is a structured screening of fifty European AI companies across four sectors (HealthTech / FinTech / HRTech-EdTech / B2B SaaS-Industrial) and fifteen-plus countries. The report applies this card to the fifty firms and produces three categories of finding: a sector-level distribution of classifications, five anonymised critical cases, and a deployer-cascade observation in §3.1 (the "Deployer Multiplier" finding). The report cites this card by edition (April 2026 v1.1) and regulatory-freeze (March 2026); a reviewer can reproduce any single classification using this card, the published primary sources, and the disclosure level documented in the parent's §6 sector tables.

2.2 Annex A — The Deployer Multiplier (May 2026)

The annex (Sprinkling Act, May 6, 2026; Zenodo DOI 10.5281/zenodo.20042175, license CC BY 4.0) develops §3.1 of the parent for three downstream segments: HRTech deployers (Annex III §4 trigger), Healthcare deployers (Article 6(1) MDR/IVDR pathway), and Retail-banking deployers (Annex III §5(b) trigger). The annex names, for each segment, the active Article 26 paragraph subset, the Article 27 FRIA logic, the Article 25 reverse-bascule check, and the trigger types that surface the obligation in operational practice. The annex is the deployer-side counterpart to this card's gate logic; it does not replace this card and does not introduce additional gates. Where this card identifies a high-risk classification under Gate 02 or Gate 03, the annex describes the deployer-side documentary cascade that follows.

2.3 What this card adds beyond the two reports

The two reports apply the framework to specific samples. This card is the framework itself. A reviewer who reads only the parent report sees the findings; a reviewer who reads only the annex sees the deployer cascade; a reviewer who reads this card sees the methodological mechanism. The three artefacts are mutually independent in their immediate object (sample, cascade, framework) and mutually consistent in their internal logic (the same six gates, the same article-mapping discipline, the same verification protocol).

2.4 What this card deliberately does not develop

This card describes the framework. It does not describe the operational heuristics by which a particular use case is mapped to a particular gate (the framework's value resides in part in those heuristics, which are exercised case by case rather than codified as a checklist). It does not describe the qualitative evaluation that allows a stability indicator to be assigned to a borderline classification (that evaluation is a judgement, not a rule). It does not describe the case-specific adjustments that follow from the contextual particulars an organisation provides at intake. The card is intended to be sufficient for a third-party reviewer to verify the conclusions of the parent report and the annex; it is not intended to be sufficient for a reader to reproduce a personalised assessment from scratch.

3. The six regulatory gates

The six gates are the operational core of the framework. They evaluate, in sequence, whether the AI system in scope triggers the most consequential regulatory regimes of the AI Act. The sequence is not arbitrary; it reflects the logical structure of Title II (prohibited practices), Title III (high-risk systems), and Title IV (general-purpose AI models and transparency obligations).

3.1 Summary table

Gate	Article	Trigger	Outcome on trigger
01	Art. 5	One of the eight prohibited AI practices: subliminal manipulation, exploitation of vulnerabilities, social scoring, predictive criminal profiling, untargeted facial-image scraping, emotion recognition in workplace and education, biometric categorisation by sensitive attribute, real-time remote biometric identification in public spaces (subject to narrow law-enforcement exceptions).	**PROHIBITED**. The system cannot be deployed legally. Assessment ends.
02	Art. 6(1)	The AI is a safety component of a product covered by Annex I Union harmonisation legislation (notably MDR Regulation (EU) 2017/745 and IVDR Regulation (EU) 2017/746) *and* the product is required to undergo third-party conformity assessment by a Notified Body.	**HIGH RISK**. Articles 8 to 22 obligations apply. Conformity assessment is integrated with the underlying sectoral assessment per Article 11(2) and Recital 81.
03	Art. 6(2) + Annex III	The AI use case falls within one of the eight Annex III categories: biometrics; critical infrastructure; education and vocational training; employment, workers management and access to self-employment; access to and enjoyment of essential private services and public services and benefits; law enforcement; migration, asylum and border control management; administration of justice and democratic processes.	**HIGH RISK**, subject to the Article 6(3) narrow exception (where the AI does not pose a significant risk of harm and meets one of four cumulative conditions).
04	Art. 51 + 55	The system embeds, or itself constitutes, a general-purpose AI model with systemic risk: training compute exceeds 10^{25} FLOPs (Article 51(1)(a)), or the model is designated by the Commission under Article 51(1)(b) on the basis of further criteria.	**HIGH RISK + GPAI systemic-risk regime**: Article 55 obligations on adversarial testing, model evaluation, incident reporting, and cybersecurity.
05	Art. 50	The AI directly interacts with natural persons (chatbots, voice assistants), generates synthetic audio / image / video / text content (deepfakes, generated text presented as informational), recognises emotions, or categorises persons biometrically — without already triggering Gate 01-04.	**LIMITED RISK**. Disclosure and labelling obligations of Article 50 §§1-5 apply; no high-risk regime is triggered on Article 50 alone.
06	Art. 53	The system embeds, or itself constitutes, a general-purpose AI model placed on the EU market. Article 53 applies to all GPAI providers; in the gate sequence, Gate 06 captures providers not already absorbed by Gate 04 (systemic-risk GPAI).	**GPAI standard obligations**: Article 53 technical documentation, training-data summary, copyright compliance, instructions for use to downstream providers.

Each gate is detailed in §§3.2 to 3.7 with its trigger condition, outcome, article cross-references, and stability indicator.

3.2 Gate 01 — Article 5 prohibited AI practices

STABLE

Article 5 of Regulation (EU) 2024/1689 lists the AI practices that are prohibited within the Union. The eight practices are not graduated: a system that triggers any one of them cannot be deployed legally, regardless of any further classification at lower gates. The Article 5 prohibitions have been applicable since February 2, 2025 (Article 113(a)), making them the earliest enforceable layer of the Act.

The eight prohibited practices

Article 5(1) lists eight categories: **(a)** AI systems that deploy subliminal, manipulative, or deceptive techniques materially distorting behaviour and causing significant harm; **(b)** AI systems exploiting vulnerabilities related to age, disability, or socio-economic situation; **(c)** social-scoring systems by public authorities or on their behalf, where the score leads to detrimental treatment unrelated to or disproportionate to the original context; **(d)** AI systems making risk assessments of natural persons solely on profiling or personality traits to predict criminal offences (with a narrow carve-out for AI systems supporting human assessment based on objective verifiable facts); **(e)** AI systems creating or expanding facial-recognition databases through untargeted scraping of facial images from the internet or CCTV; **(f)** AI systems inferring emotions in workplaces and educational institutions, except for medical or safety reasons; **(g)** biometric categorisation systems individually categorising natural persons on the basis of their biometric data to deduce or infer race, political opinions, trade-union membership, religious or philosophical beliefs, sex life, or sexual orientation; **(h)** real-time remote biometric identification in publicly accessible spaces for law-enforcement purposes, subject to the strict exceptions of Article 5(2) to (7).

Outcome and consequences

A trigger at Gate 01 ends the assessment. The system cannot be deployed legally in the EU. The classification is dated and recorded but no further gate is evaluated. The penalty regime of Article 99(3) for breaches of Article 5 is the most severe in the Act: up to EUR 35 million, or 7% of total worldwide annual turnover for the preceding financial year, whichever is higher.

Stability indicator

STABLE for most categories. The text of the eight practices is precise. The exceptions in points (d), (f), and (h) are narrow and have been further specified by Commission guidelines on prohibited AI practices (April 2025). Cases that turn on the boundary between point (a) "subliminal techniques" and ordinary marketing AI may reach MODERATE pending further enforcement clarification. Real-time biometric ID in public spaces under (h) is STABLE on its prohibition but national implementing legislation will determine the operational exception architecture per Article 5(5).

3.3 Gate 02 — Article 6(1) safety-component pathway

STABLE

Article 6(1) of the AI Act establishes that an AI system is high-risk where two cumulative conditions are met: (a) the AI system is intended to be used as a safety component of a product covered by the Union harmonisation legislation listed in Annex I, or is itself such a product; *and* (b) the product whose safety component is the AI system, or the AI system itself, is required to undergo a third-party conformity assessment, with a view to the placing on the market or the putting into service of that product, pursuant to that legislation.

Annex I list

Annex I, Section A lists the existing Union harmonisation legislation under which Article 6(1) applies. The list includes Regulation (EU) 2017/745 on medical devices (MDR), Regulation (EU) 2017/746 on in-vitro diagnostic medical devices (IVDR), Regulation (EU) 2018/858 on motor vehicles, Regulation (EU) 2017/746 on machinery (Machinery Regulation), Regulation (EU) 2018/1139 on civil aviation, Directive 2014/90/EU on marine equipment, and several others. Annex I, Section B lists further legislation that is not currently subject to AIA Title III obligations (notably toys, lifts, recreational craft, pressure equipment, gas appliances) but where the AIA framework foresees a future extension.

Medical AI specificity

For AI systems qualified as Medical Device Software (MDSW) under MDR / IVDR, the second condition (third-party conformity assessment) depends on the 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 page 6.

Outcome and integration

A trigger at Gate 02 classifies the system as high-risk. Articles 8 to 22 of the AI Act apply: risk-management system (Art. 9), data governance (Art. 10), technical documentation (Art. 11), record-keeping / automatic logging (Art. 12), transparency and provision of information to deployers (Art. 13), human oversight (Art. 14), accuracy, robustness and cybersecurity (Art. 15), and the obligations of Articles 16 to 22 on providers. Article 11(2) of the AI Act and Recital 81 require integration of the AI Act technical documentation with the existing sectoral conformity-assessment file (e.g. the MDR technical documentation file) rather than duplication. The Notified Body that assesses MDR / IVDR conformity is the same body that assesses AI Act conformity for the Article 6(1) pathway, where the Notified Body is designated under both regimes.

Stability indicator

STABLE for the basic two-condition test. The MDR-IVDR-AIA mapping in MDCG 2025-6 is Commission-endorsed guidance and is operationally settled. Cases that turn on whether a given AI use is a "safety component" within Article 3(14) AIA may reach MODERATE pending further sector-specific guidance.

3.4 Gate 03 — Article 6(2) plus Annex III high-risk use cases MODERATE

Article 6(2) of the AI Act establishes that AI systems referred to in Annex III are high-risk. Annex III lists eight categories of use case in which AI systems are considered high-risk irrespective of whether they are safety components of regulated products. The Annex III gate is the operational entry point for high-risk classification of standalone AI systems (i.e. systems not covered by Gate 02). It is the gate most frequently triggered in operational practice.

The eight Annex III categories

Annex III lists: **§1** biometrics (remote biometric identification, biometric categorisation, emotion recognition outside the Article 5 prohibitions); **§2** critical infrastructure (safety components in critical digital infrastructure, road traffic, water, gas, heating, electricity supply); **§3** education and vocational training (admission, evaluation, allocation between institutions, monitoring of prohibited behaviour during tests); **§4** employment, workers management and access to self-employment (recruitment, screening, evaluation, promotion, allocation of tasks based on individual behaviour or personal traits); **§5** access to and enjoyment of essential private services and essential public services and benefits, divided into (a) public benefits eligibility evaluation, (b) creditworthiness evaluation and credit-scoring, (c) risk assessment and pricing for life and health insurance, (d) emergency dispatch and patient-triage; **§6** law enforcement (specific subcategories for risk assessment, polygraphs, deepfake detection, evidence reliability, profiling); **§7** migration, asylum and border-control management (polygraphs, risk assessment, application examination, recognition); **§8** administration of justice and democratic processes (assistance to judicial authorities in researching and interpreting facts and law, influencing electoral outcomes).

Article 6(3) narrow exception

Article 6(3) introduces a narrow exception: an AI system referred to in Annex III is not high-risk where it does not pose a significant risk of harm to health, safety, or fundamental rights, and one of four cumulative conditions is met: (a) the AI is intended to perform a narrow procedural task; (b) the AI is intended to improve the result of a previously completed human activity; (c) the AI is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment without proper human review; (d) the AI is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases listed in Annex III. The exception does not apply where the AI system performs profiling of natural persons. The framework applies the exception conservatively: where doubt exists, the Article 6(3) exception is not granted, and the classification remains HIGH RISK with the recommendation to seek qualified legal counsel.

Outcome

A trigger at Gate 03 classifies the system as HIGH RISK under Article 6(2). The full set of Articles 8 to 22 obligations applies. The deployer-side obligations of Article 26 and (where applicable) Article 27 FRIA apply additionally. The annex registration of Article 49 applies for public-law-body deployers. *Annex A* develops the deployer-side cascade for the three most operationally consequential Annex III triggers (§4 employment, §5(b) credit-scoring, plus the Article 6(1) MDR/IVDR pathway for healthcare, which is not technically an Annex III trigger but functions as the practical equivalent for medical AI).

Stability indicator

MODERATE for several Annex III categories. The AI Office has signalled progressive guidance on §1 biometrics, §5 access to services, and §6 law enforcement, where the boundary between high-risk and the Article 6(3) exception remains under interpretive evolution. STABLE for §4 employment and §5(b) credit-scoring, where the operational scope is largely settled and where binding case law (notably SCHUFA C-634/21 for §5(b)) reinforces the high-risk reading.

3.5 Gate 04 — Article 51 plus 55 GPAI with systemic risk

UNSTABLE

Article 51 of the AI Act establishes the classification of general-purpose AI models with systemic risk. A GPAI model is classified as having systemic risk where it has high-impact capabilities; capabilities are presumed high-impact under Article 51(2) where the cumulative amount of compute used for training exceeds 10^{25} floating-point operations (FLOPs). The Commission may also designate a model as having systemic risk on the basis of further criteria listed in Annex XIII (parameters, dataset size, energy consumption, market reach, modalities, registered users).

Trigger condition

The framework triggers Gate 04 where the AI system in scope embeds a model that has been classified as having systemic risk under Article 51 (whether by compute presumption or by Commission designation), or where the system itself is such a model. For deployers integrating a third-party GPAI model, the trigger is the inheritance of obligations through Article 13 instructions for use; for providers of the GPAI model itself, the trigger is the direct application of Article 55.

Article 55 obligations

Providers of GPAI models with systemic risk must, in addition to the standard Article 53 obligations: **(a)** perform model evaluation in accordance with standardised protocols and tools reflecting the state of the art, including conducting and documenting adversarial testing of the model with a view to identifying and mitigating systemic risks; **(b)** assess and mitigate possible systemic risks at Union level, including their sources, that may stem from the development, the placing on the market, or the use of the model; **(c)** keep track of, document, and report serious incidents and possible corrective measures to the AI Office and, as appropriate, to national competent authorities; **(d)** ensure an adequate level of cybersecurity protection for the model and the physical infrastructure of the model.

Outcome

A trigger at Gate 04 classifies the system as HIGH RISK plus the GPAI systemic-risk regime. The Article 55 obligations apply in addition to the Article 53 standard GPAI obligations and, where the AI system itself is a high-risk system under Gate 02 or Gate 03, the Articles 8 to 22 obligations apply concurrently. The penalty regime of Article 101 specifically addresses GPAI providers and is graduated.

Stability indicator

UNSTABLE on the boundary. The 10^{25} FLOPs threshold is operationally measurable but the boundary as to which models cross it is evolving as a function of Commission designations and as a function of provider self-disclosure. The criteria of Annex XIII for further designation are not yet fully operationalised in published Commission decisions; the AI Office is progressively building the decisional record. STABLE on the consequence: where a model is recognised as having systemic risk, Article 55 applies in full.

3.6 Gate 05 — Article 50 transparency

STABLE

Article 50 of the AI Act introduces transparency and labelling obligations on certain categories of AI systems and certain categories of AI-generated content. The obligations apply irrespective of whether the system has triggered Gates 01 to 04; they are additional rather than substitutive. Article 50 becomes binding on August 2, 2026 — the general application date under Article 113.

The four sub-obligations

Article 50(1): providers of AI systems intended to interact directly with natural persons must design and develop them in such a way that the natural persons concerned are informed they are interacting with an AI system, unless this is obvious from the circumstances or use. **Article 50(2)**: providers of AI systems generating synthetic audio, image, video, or text content (including general-purpose AI systems generating such content) must ensure the outputs are marked in a machine-readable format and detectable as artificially generated or manipulated. **Article 50(3)**: deployers of emotion-recognition or biometric-categorisation systems (where these are not prohibited under Article 5) must inform the natural persons exposed thereto of the operation of the system and process their personal data in accordance with the GDPR and Regulation (EU) 2018/1725. **Article 50(4)**: deployers of AI systems generating or manipulating image, audio, or video content constituting a deepfake must disclose that the content has been artificially generated or manipulated, with limited carve-outs for evident artistic, creative, satirical, fictional, or analogous works.

Outcome

A trigger at Gate 05 classifies the system as LIMITED RISK. Disclosure obligations apply in the form, modality, and timing specified in Article 50; no high-risk obligations are triggered on Article 50 alone. Gate 05 applies as an overlay on top of any existing classification: a HIGH RISK system under Gates 02 or 03 that also constitutes an Article 50(1) chatbot must comply with both regimes.

Stability indicator

STABLE for the four sub-obligations. MODERATE on the operational format of "machine-readable marking" required by Article 50(2), where standardised marking protocols (the AI Office is supporting the development of harmonised standards on watermarking and provenance signalling) have not yet stabilised. STABLE on the deepfake disclosure obligation of Article 50(4) and on its narrow carve-outs.

3.7 Gate 06 — Article 53 GPAI standard obligations

STABLE

Article 53 establishes the obligations applicable to all providers of general-purpose AI models placed on the EU market, regardless of whether the model has systemic risk under Article 51. The four obligations are: **(a)** draw up and keep up-to-date the technical documentation of the model, including its training and testing process and the results of its evaluation, with the elements set out in Annex XI; **(b)** draw up, keep up-to-date and make available information and documentation to providers of AI systems who intend to integrate the GPAI model into their AI systems, with the elements set out in Annex XII; **(c)** put in place a policy to comply with Union law on copyright and related rights, in particular to identify and comply with reservations of rights expressed under Article 4(3) of Directive (EU) 2019/790; **(d)** draw up and make publicly available a sufficiently detailed summary of the content used for training of the model, according to a template provided by the AI Office.

A trigger at Gate 06 classifies the model under the Article 53 standard regime. The obligations apply from August 2, 2025 for new models and from August 2, 2027 for models placed on the market before August 2, 2025 (Article 111(3)). Stability is STABLE on the four obligations; MODERATE on the operational implementation of the AI Office training-data summary template, which is in active development.

4. Article-level mapping — bias, fairness, and human oversight

Once a system is classified as high-risk under Gate 02 or Gate 03, the substantive obligations of Articles 8 to 15 apply. Four of those articles carry the analytical weight of the bias-and-fairness frontier: Article 9 (risk-management system), Article 10 (data and data governance), Article 14 (human oversight), and Article 15 (accuracy, robustness, and cybersecurity). This section maps the framework to those four articles with the precision required for a third-party reviewer with academic background in algorithmic fairness to evaluate whether the framework treats the bias dimension seriously.

4.1 Article 9 — risk-management system

Article 9 requires that a risk-management system be established, implemented, documented, and maintained as a continuous iterative process throughout the entire lifecycle of the high-risk AI system. The article requires identification and analysis of known and reasonably foreseeable risks; estimation and evaluation of risks that may emerge under intended use and reasonably foreseeable misuse; evaluation of risks based on the analysis of post-market-monitoring data (Article 72); and adoption of appropriate and targeted risk-management measures. Within the framework, Article 9 evidence is documented at the gate-output level: each high-risk classification carries a risk-management documentation pointer indicating whether the deployer or provider has constituted, in operational form, the risk-management system the article requires.

4.2 Article 10 — data and data governance

Article 10 carries the most analytically dense provisions on bias and fairness. **Article 10(2)(f)** requires examination in view of possible biases that are likely to affect the health and safety of natural persons, have a negative impact on fundamental rights, or lead to discrimination prohibited under Union law. **Article 10(2)(g)** requires appropriate measures to detect, prevent, and mitigate such biases. **Article 10(3)** requires that training, validation, and testing datasets be relevant, sufficiently representative, and to the best extent possible free of errors and complete in view of the intended purpose. **Article 10(5)** permits, on the basis of strict conditions, the processing of special categories of personal data within the meaning of GDPR Article 9 for the purpose of bias detection and correction — the processing must be strictly necessary, the data must be subject to appropriate safeguards, and the dataset must be deleted once the purpose has been achieved.

Within the framework, Article 10 compliance is evaluated at the data-governance attestation level. The framework does not produce a continuous "fairness score"; it produces a binary attestation per gate (triggered / not triggered / borderline) with explicit documentation of the bias-examination evidence on file. The framework is consistent with the academic finding that aggregate fairness metrics can mask intersectional harm: the attestation does not certify "fairness" in the substantive sense; it certifies that the Article 10(2)(f) and (g) examination has been conducted and documented at a level a Notified Body or supervisory authority would accept on inspection. Substantive fairness is, in the framework's view, the responsibility of the system's designers and of the bias-detection methodology they adopt; the framework attests procedural compliance, not outcome equity.

4.3 Article 14 — human oversight

Article 14 requires that high-risk AI systems be designed and developed in such a way, including with appropriate human-machine interface tools, that they can be effectively overseen by natural persons during the period of use. Article 14(4) identifies five oversight capabilities: properly understanding the relevant capacities and limitations of the system and being able to monitor its operation; remaining aware of the possible tendency to over-rely on the output (automation bias); correctly interpreting the high-risk system's output, taking into account the available interpretation tools and methods; deciding, in any particular situation, not to use the high-risk system or otherwise to disregard, override, or reverse its output; and, where necessary, intervening or interrupting its operation. The framework documents Article 14 compliance at the operational-architecture level: an HR-AI deployer that operates with caseload caps making meaningful oversight implausible (e.g. fifty AI-ranked candidate files reviewed per hour) is documented as failing the §2 oversight test even where the formal oversight policy exists on paper.

4.4 Article 15 — accuracy, robustness, and cybersecurity

Article 15 requires that high-risk AI systems be designed and developed in such a way that they achieve, in light of their intended purpose, an appropriate level of accuracy, robustness, and cybersecurity, and that they perform consistently in those respects throughout their lifecycle. The article requires that the levels of accuracy and the relevant accuracy metrics be declared in the accompanying instructions of use. Within the framework, Article 15 evidence is documented at the technical-documentation level: each high-risk classification points to the existence of declared accuracy metrics in the technical-documentation file (Annex IV) and to the existence of the consistent-performance evidence Article 15(4) requires.

5. Annex III decomposition — eight categories

Annex III is the operational vocabulary of Gate 03. The eight categories are not equally clear-cut; some have stable boundaries, others are under active interpretive evolution by the AI Office. This section names each category, its scope, and the typical stability indicator the framework assigns at Gate 03.

§	Category	Scope (summary)	Stability
§1	Biometrics	Remote biometric identification systems (excluding biometric verification authentication); biometric categorisation systems by sensitive or protected attributes; emotion-recognition systems outside the Article 5(1)(f) workplace and education prohibition.	MODERATE
§2	Critical infrastructure	AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, and the supply of water, gas, heating, and electricity.	STABLE
§3	Education and vocational training	AI for determination of access, admission, or assignment to educational institutions; for evaluation of learning outcomes; for assessing the appropriate level of education an individual will receive or be able to access; and for monitoring and detecting prohibited behaviour during tests.	MODERATE
§4	Employment, workers management	AI for recruitment, screening, filtering, and evaluation of applications; for promotion, termination, and task allocation based on individual behaviour or personal traits; and for monitoring and evaluating performance and behaviour of employed persons.	STABLE
§5(a)	Public benefits eligibility	AI used by public authorities or on their behalf to evaluate eligibility for essential public-assistance benefits and services, including healthcare services, or to grant, reduce, revoke, or reclaim such benefits and services.	STABLE
§5(b)	Creditworthiness, credit-scoring	AI for evaluating creditworthiness of natural persons or establishing their credit score, with the exception of AI used to detect financial fraud. Reinforced operationally by CJEU SCHUFA C-634/21 (December 7, 2023) on GDPR Article 22.	STABLE
§5(c)	Insurance pricing	AI for risk assessment and pricing in life and health insurance.	MODERATE
§5(d)	Emergency dispatch and triage	AI for evaluating and classifying emergency calls by natural persons and dispatching emergency services or establishing priority for emergency healthcare patients.	STABLE
§6	Law enforcement	Six subcategories — risk assessment of natural persons; polygraph and similar tools; deepfake and evidence-reliability assessment; profiling under Directive (EU) 2016/680; investigative analysis of natural persons.	MODERATE
§7	Migration, asylum, borders	Polygraphs and similar tools; risk assessment of natural persons; AI assistance to examination of applications for asylum, visa, residence permits; AI for recognition of natural persons in border-control contexts.	MODERATE
§8	Justice, democratic processes	AI for assistance to judicial authorities in researching and interpreting facts and law and applying the law to a concrete set of facts (alternative dispute resolution covered); AI for influencing the outcome of an election or referendum or the voting behaviour of natural persons.	MODERATE

The categories most operationally consequential in current European AI deployment are §4 (employment), §5(b) (credit-scoring), and the Article 6(1) MDR/IVDR pathway for medical AI (which is not technically Annex III but functions similarly). Those three are the segments developed in *Annex A — The Deployer Multiplier* (May 2026).

6. General-purpose AI obligations regime

The general-purpose AI (GPAI) regime applies to providers of GPAI models placed on the EU market. The regime is structured around four articles: Article 50 (transparency), Article 51 (classification of GPAI with systemic risk), Article 53 (standard obligations for all GPAI providers), and Article 55 (additional obligations for systemic-risk GPAI providers). The regime applies independently of whether the GPAI model is integrated into a high-risk AI system under Articles 6 to 22.

6.1 The GPAI definition

Article 3(63) defines a general-purpose AI model as an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications. The definition excludes AI models that are used for research, development, and prototyping activities before they are placed on the market.

6.2 Provider versus deployer of a GPAI

The GPAI obligations apply primarily to the provider of the GPAI model. A deployer that integrates a third-party GPAI model into its own AI system inherits part of the regulatory burden via the Article 13 instructions for use that the GPAI provider must furnish (Article 53(1)(b) requires the GPAI provider to make available the information and documentation set out in Annex XII to providers of AI systems integrating the model). The line between provider and deployer can shift: a deployer who fine-tunes a GPAI model on its own data in a way that constitutes a substantial modification within the meaning of Article 3(23) becomes a provider of the modified model under Article 25(1)(b) and inherits the corresponding obligations.

6.3 Article 53 standard obligations

All GPAI providers (irrespective of systemic-risk status) must comply with Article 53: technical documentation per Annex XI; information and documentation to downstream providers per Annex XII; copyright-compliance policy per Article 4(3) of Directive (EU) 2019/790; publicly available training-data summary per a template to be issued by the AI Office.

6.4 Article 55 additional obligations for GPAI with systemic risk

GPAI providers whose models meet the Article 51 threshold (presumed at $>10^{25}$ FLOPs of training compute, or designated by the Commission) must additionally: perform model evaluation including adversarial testing; assess and mitigate systemic risks at Union level; track, document, and report serious incidents to the AI Office; ensure cybersecurity protection of the model and the supporting infrastructure.

6.5 Operational positioning of the GPAI regime

The framework treats Gate 04 (Article 51 + 55) as primary for systemic-risk GPAI providers and Gate 06 (Article 53) as primary for non-systemic-risk GPAI providers. For deployers integrating a GPAI model, the framework documents which obligations of the upstream GPAI provider have been formally received (in particular the Article 53(1)(b) Annex XII documentation), and which residual deployer-side obligations apply (notably Article 50 transparency for any user-facing AI system embedding the model, and Article 26 deployer obligations where the embedding system is itself high-risk).

7. Deployer-side obligations and the cascade

Article 26 of the AI Act sets out the obligations of deployers of high-risk AI systems. Article 27 sets out the Fundamental Rights Impact Assessment (FRIA) obligation that applies to certain categories of deployers. Together with Articles 49 (registration in the EU database for public-law-body deployers) and 50 (transparency obligations that apply also to certain deployers), they form the deployer-side regulatory layer of the AI Act. The framework documents this layer in summary form here; the operational cascade across three downstream segments is developed in *Annex A — The Deployer Multiplier* (May 2026).

7.1 Article 26 — twelve paragraphs

Article 26 sets out twelve obligations on deployers: **§1** use the system in accordance with the instructions for use and take appropriate technical and organisational measures; **§2** assign human oversight to natural persons who have the necessary competence, training, authority, and support; **§3** ensure that input data is relevant and sufficiently representative in view of the intended purpose; **§4** monitor the operation of the high-risk AI system on the basis of the instructions for use; **§5** keep the logs automatically generated by the system; **§6** retain those logs for a period appropriate to the intended purpose, and at least six months; **§7** for deployers in the workplace, inform workers' representatives and affected workers before putting the system into service; **§8** for public-authority deployers, register the use in the EU database referred to in Article 49; **§9** use the information provided under Article 13 to comply with the obligation to carry out a data protection impact assessment under GDPR Article 35; **§10** in cases of biometric identification, follow specified post-remote-biometric-identification procedures; **§11** inform natural persons subject to the use of a high-risk AI system in Annex III categories; **§12** cooperate with competent authorities on any action they take in respect of the high-risk AI system. The active subset varies by deployment context; *Annex A* documents the active subset for HRTech deployers (§§2, 5, 6, 7, 11), Healthcare deployers (§§1, 2, 5, 6, 9, 12), and Retail-banking deployers (variable subset including FRIA-related §3, §5, §11).

7.2 Article 27 — FRIA

Article 27 requires deployers that are bodies governed by public law, or private entities providing public services, deploying a high-risk Annex III system (with carve-outs for the §2 critical-infrastructure category) to perform, before putting the system into service, an assessment of the impact on fundamental rights that the use of the system may produce. The FRIA covers the use process, the period of intended use, the categories of natural persons likely to be affected, the specific risks of harm, the measures of human oversight, and the measures to take in case the risks materialise. The deployer notifies the market-surveillance authority of the assessment conclusions. Article 27 applies to Article 6(2) Annex III systems only; Article 6(1) safety-component systems are not within Article 27 scope (the framework treats this boundary as STABLE; *Annex A* §5.5 develops the consequences for healthcare deployers operating under the MDR/IVDR pathway).

7.3 Article 25(1)(b) reverse-bascule

Article 25(1)(b) classifies as a provider any person who substantially modifies a high-risk AI system already placed on the market or put into service. Once classified as a provider under Article 25(1)(b), the entity inherits the heavier provider regime of Articles 16 to 22 (technical documentation, conformity assessment, post-market monitoring, quality management system). The framework includes a reverse-bascule check for any deployment where fine-tuning, retraining, or substantial parameter modification of the vendor model occurs on the deployer's data (workforce data, patient conversations, credit-decision history). The check is systematic in the framework, even when the answer is "no modification beyond intended purpose"; the documentation of the negative finding is part of the audit trail.

7.4 Pointer to Annex A for the operational cascade

The deployer-side obligations are developed operationally in *Annex A — The Deployer Multiplier* (Sprinkling Act, May 6, 2026; Zenodo DOI 10.5281/zenodo.20042175). The annex documents, for three downstream segments (HRTech, Healthcare, Retail-banking): the active Article 26 paragraph subset, the Article 27 FRIA logic, the Article 25 reverse-bascule worked example for healthcare, and the four trigger types (incident, supervisor, contract, internal audit) under which the obligation surfaces in operational practice. This card and the annex are mutually consistent: the gates this card defines are the gates the annex applies; the article mapping this card establishes is the article mapping the annex develops onto the deployer's seat.

8. Output of an assessment — the artefact format

Every assessment produced under this framework results in a structured artefact. The artefact has six standard components, in fixed order, so that two assessments produced at different dates and on different cases are immediately comparable and so that a third-party reviewer can locate the same information at the same place in any assessment.

8.1 Identification block

Subject of assessment, identifier of the AI system or use case, deployer or provider in scope, intake date, methodology edition (this card, e.g. April 2026 v1.1), regulatory freeze date, badge code (SA-YYYYMMDD-NNNN format).

8.2 Gate-by-gate path

For each of the six gates, in sequence: trigger condition tested; result (triggered / not triggered / borderline); article cross-reference; rationale; stability indicator (STABLE / MODERATE / UNSTABLE). The path terminates at the first gate triggered (Gate 01) or, where Gate 01 is not triggered, runs through all six gates with a documented outcome at each.

8.3 Final classification

One of: PROHIBITED (Gate 01); HIGH RISK (Gate 02 or 03 or 04); LIMITED RISK (Gate 05 only); GPAI STANDARD (Gate 06 only); UNCLASSIFIED (the system falls outside the AI Act scope). The final classification is the gate result, not a synthetic score.

8.4 Active obligations subset

Where the classification is HIGH RISK, the active subset of Articles 8 to 22 (provider) or Article 26 (deployer) is named explicitly. Where Article 27 FRIA applies, the FRIA scope is named. Where Article 25(1)(b) reverse-bascule check is applicable, the result of the check is documented even where the result is negative.

8.5 Stability statement

The dominant stability indicator across the path is named (the lowest stability across the gates that contributed to the final classification). The classification's expected lifespan is documented: for UNSTABLE classifications, a re-assessment is recommended within twelve months or upon any material regulatory development; for STABLE classifications, the next-due re-assessment is set against the next regulatory-freeze date.

8.6 Verification block

SHA-256 of the published PDF; OpenTimestamps proof identifier; mirror URL on github.com/sprinkling-act/timestamps; badge URL on sprinklingact.com/verify/{badge-code}; recommended citation block (suggested APA or BibTeX). The verification block is what makes the artefact independently auditable: a third party can verify the date and integrity without contacting Sprinkling Act.

9. Verification and versioning protocol

9.1 Cryptographic timestamp

Every Sprinkling Act published artefact (this card, the parent screening report, the annex) is timestamped via OpenTimestamps and anchored on the Bitcoin blockchain through four independent calendar servers (a/b OpenTimestamps pools, Eternity Wall, Catallaxy). The detached proof file (`.ots`) is published alongside the PDF and mirrored on the public repository github.com/sprinkling-act/timestamps. The proof binds the SHA-256 of the published PDF to a Bitcoin block confirmation, allowing any third party to verify the date of the document independently of Sprinkling Act.

9.2 Verification procedure

To verify a published artefact: **(1)** install the OpenTimestamps client (`pip3 install opentimestamps-client`); **(2)** download the PDF and the matching `.ots` file from the published source (sprinklingact.com or the GitHub mirror or, where applicable, Zenodo); **(3)** run `ots verify {filename}.pdf.ots` — the client returns the calendar attestations and the Bitcoin block confirmation. Where the document has not yet been upgraded with a Bitcoin block (the upgrade typically occurs within one hour of stamping), the calendars provide pending attestations that can be re-verified after upgrade.

9.3 Methodology versioning

The methodology card is versioned by edition (publication month) and sub-version (point releases within an edition, e.g. v1.1). The framework's commitment is that any artefact published under a given edition is reproducible against that exact version of the methodology. Version archives are retained on Zenodo with concept-DOI resolution to the latest version and version-specific DOI for historical reproducibility.

9.4 Regulatory freeze date

The regulatory freeze date is the date through which regulatory developments have been incorporated into the methodology. The framework's commitment is that any classification produced under a given regulatory freeze date is grounded in the regulatory state of that date, not in subsequent developments. A classification produced under a March 2026 regulatory freeze that no longer holds under a subsequent freeze (e.g. because the Digital Omnibus on AI is adopted and shifts a deadline) is documented as superseded but not retroactively amended; the published artefact carries its freeze date with it.

9.5 Errata and corrections

Material errors discovered after publication are corrected through versioned re-publication (v1.1, v1.2, …) with an explicit change log. The OpenTimestamps proof of the corrected version is generated; the original proof remains valid for the original version. Where the correction reverses a published classification, the affected client is notified directly and the published artefact is annotated with the supersession reference.

10. Alignment with international risk-management standards

The framework is structurally consistent (not formally certified) with three international frameworks. **NIST AI Risk Management Framework 1.0** (NIST AI 100-1, 2023): the gate-by-gate evaluation maps to the Map and Measure functions of the NIST framework. **ISO/IEC 42001:2023** (AI Management System): the six-gate output produces the risk classification and obligation mapping that feeds into an ISO 42001-compliant AI Management System. **WHO Regulatory considerations on AI for health** (2023, ISBN 9789240078871) and **WHO Generating evidence for AI-based medical devices** (2021, ISBN 9789240038462): the six topic areas of the WHO regulatory considerations are mapped to AI Act and MDR articles in healthcare assessments. Sprinkling Act does not claim NIST AI RMF compliance, ISO 42001 certification, or WHO endorsement; the references indicate structural coherence only.

11. Limitations and qualified legal counsel

11.1 Interpretive limits

The EU AI Act requires contextual interpretation. The framework does not resolve interpretive ambiguity; it flags ambiguity explicitly and recommends qualified legal counsel for borderline classifications. Where two readings of a provision are equally defensible, the framework names both rather than selecting the lower-risk reading. The user of the framework's output is responsible for selecting between the named readings or for instructing legal counsel to do so.

11.2 Operational, not legal

The framework is an operational tool. It produces the structured, article-mapped artefact a lawyer, regulator, or investor needs as a starting point. It does not produce legal advice, a binding legal opinion, or a determination of compliance within the meaning of any provision of the AI Act. A classification produced under the framework can be input into a legal proceeding or a conformity-assessment file; it cannot stand alone as the determination on which a final compliance decision rests.

11.3 Input dependence

The classification produced by the framework rests on the inputs provided by the deployer or provider. Incomplete or inaccurate inputs produce incomplete or inaccurate classifications. The framework does not validate the truth of the deployer's or provider's representations about the AI system; it documents the classification implied by those representations and flags inconsistencies where they appear within the input itself.

11.4 Scope limits

The framework does not cover: national implementing legislation (the AI Act allows substantial member-state-level variation in implementation, particularly on Article 5 exceptions, Article 99 sanctions, and the architecture of national supervisory authorities); sector-specific regulations beyond those expressly cross-referenced in the AI Act (notably MDR, IVDR, and Annex I product legislation); the GDPR beyond points where the AI Act explicitly cross-references it; and AI systems deployed outside the EU. Cross-border deployments may engage additional regulatory regimes that fall outside the framework.

11.5 Classification is not certification

A favourable classification under this framework is not a certificate of compliance. The framework's output is not a substitute for the conformity assessment that Article 43 requires for certain high-risk AI systems, for the Notified Body assessment that the corresponding sectoral legislation may require, or for the registration with the AI Office that Article 49 requires for public-authority deployers. The framework's value is upstream of those processes; it documents the classification on which those processes will operate, not the processes themselves.

11.6 Independence and conflicts

Sprinkling Act is an independent advisory firm. It is not a law firm, not a Notified Body, not a certification body, and is not affiliated with, endorsed by, or acting on behalf of the European Commission, the European Parliament, the AI Office, the Artificial Intelligence Board, or any national supervisory authority. Where Sprinkling Act has a commercial relationship with a firm whose system is being assessed in published research, the relationship is disclosed in the relevant publication.

12. References

Primary regulatory sources. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, OJ L of 12 July 2024 (the "AI Act"). Regulation (EU) 2017/745 on medical devices (MDR). Regulation (EU) 2017/746 on in-vitro diagnostic medical devices (IVDR). Directive (EU) 2019/790 on copyright in the Digital Single Market. Regulation (EU) 2016/679 (GDPR).

Commission and AI Office guidance. MDCG 2025-6 / AIB 2025-1, *Interplay between Medical Device Regulation, In Vitro Diagnostic Medical Device Regulation, and the Artificial Intelligence Act*, June 19, 2025. Commission Guidelines on prohibited AI practices (April 2025). Commission Implementing Regulation on the AI Office training-data summary template (in development).

Binding case law. CJEU, SCHUFA Holding AG, Case C-634/21, judgment of 7 December 2023 (on GDPR Article 22 and credit-scoring).

International standards. NIST AI Risk Management Framework 1.0 (NIST AI 100-1, January 2023). ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system. WHO, *Regulatory considerations on artificial intelligence for health*, 2023, ISBN 9789240078871. WHO, *Generating evidence for artificial intelligence-based medical devices*, 2021, ISBN 9789240038462.

Sprinkling Act published artefacts. Sprinkling Act, *EU AI Act Readiness Report*, April 12, 2026, Zenodo DOI 10.5281/zenodo.19671329. Sprinkling Act, *Annex A — The Deployer Multiplier (May 2026)*, May 6, 2026, Zenodo DOI 10.5281/zenodo.20042175.

13. Document metadata

Field	Value
Edition	April 2026 (v1.1)
Internal version code	SA-METH-2026.04 (legacy traceability identifier; not used as branding)
Regulatory freeze	March 2026
Publication date	April 27, 2026
Source of authority	Regulation (EU) 2024/1689 — OJ L of 12 July 2024
Companion documents	Parent screening report (Zenodo DOI 10.5281/zenodo.19671329); Annex A (Zenodo DOI 10.5281/zenodo.20042175)
License	Creative Commons Attribution 4.0 International (CC BY 4.0)
Author	Lamar B. Shucrani, Founder, Sprinkling Act. Brussels-based AI Act analyst working on the pre-conformity layer for European deployers and providers.
Author ORCID	0009-0002-5093-8550
Cryptographic timestamp	OpenTimestamps detached proof (.ots) anchored to Bitcoin via four calendar servers. Mirror: github.com/sprinkling-act/timestamps
Suggested citation	Shucrani, L. B. (2026). *Sprinkling Act Methodology Card — April 2026 edition*. Sprinkling Act. https://sprinklingact.com/methodology
Errata channel	Versioned re-publication with change log; OpenTimestamps proof of each edition remains valid for that edition.
Independence	Not a law firm, not a Notified Body, not a certification body, not affiliated with the European Commission, the European Parliament, the AI Office, the Artificial Intelligence Board, or any national supervisory authority.
Correspondence	contact@sprinklingact.com

REFERENCE SUMMARY

Sprinkling Act Methodology Card — April 2026 edition

This card consolidates the assessment framework in 13 sections. It documents the six-gate logic, the article-level mapping (Articles 9, 10, 14, 15), the deployer-side cascade (Articles 26, 27, 49, 50), the GPAI obligations regime (Articles 50, 51, 53, 55), the Annex III decomposition, the output artefact format, the verification protocol, and the alignment with international risk-management standards (NIST AI RMF, ISO/IEC 42001, WHO).

SUGGESTED CITATION

Shucrani, L. B. (2026). *Sprinkling Act Methodology Card — April 2026 edition*. Sprinkling Act. https://sprinklingact.com/methodology

Companion documents
Parent EU AI Act Readiness Report (April 2026)
Annex A — The Deployer Multiplier (May 2026)

Apply the framework
sprinklingact.com/check
9 questions · 60 seconds · zero data collected