



EU AI ACT BRIDGE SERIES · PRACTITIONER GUIDE 03

Is Your AI System *High-Risk*? An Annex III Classification Guide

A guided decision-tree through the EU AI Act Annex III screening process — with explanatory notes, real-world examples for each of the eight classification areas, and a completed screening record you can sign off and retain as evidence.

DOCUMENT ID	VERSION
EB00K-AIMS-CLASS-001	1.0 · September 2026

LEGAL BASIS

Regulation (EU) 2024/1689 · Art. 6 · Annex III

AUDIENCE

DPOs · CIOs · Legal & Compliance Leads · AI Owners

THIS GUIDE IS AN EDUCATIONAL REFERENCE AND DOES NOT CONSTITUTE LEGAL ADVICE. CLASSIFICATION DECISIONS WITH LEGAL CONSEQUENCES SHALL BE REVIEWED BY QUALIFIED LEGAL COUNSEL. ALL STATUTORY DEADLINES ARE CORRECT AS OF SEPTEMBER 2026 AND SHALL BE VERIFIED AGAINST CURRENT OFFICIAL SOURCES.

EU AI ACT · ART. 6 + ANNEX III

ISO/IEC 42001:2023

ITIL 4 · RISK MANAGEMENT

COMPANION: REG-AIMS-CLASS-001 EXCEL TEMPLATE

CONTENTS

What Is Inside

01	The Classification Problem	Art. 6(2)
	Why "we don't believe our tools are high-risk" is not a compliance record — and what the EU AI Act actually demands	
02	Your Legal Obligation	Art. 6(1)(2)(3)
	Article 6 and Annex III explained — the gateway clause, its three sub-provisions, and the obligation cascade that follows from a High-Risk finding	
03	Before You Begin	Scope + Inventory
	Scoping your AI inventory — what counts as an AI system, what is excluded, and why every system needs its own assessment row	
04	The Three-Step Classification Test	Steps 1-3
	The structured decision process that produces your classification record: area screening → use-case specificity → Art. 6(3) exclusion	
05	The Eight Annex III Areas — Area by Area	Areas 1-8
	Classification logic, explanatory notes, and real-world examples for each of the eight areas — with a verdict and conditional caveat where applicable	
06	Article 6(3) — The Narrow Exclusion	Art. 6(3)
	When the exclusion is available, what evidence it requires, and the most common conflation that invalidates it	
07	If You Find a High-Risk System	Arts. 9-15 · Art. 26

The obligation cascade — Arts. 9–15, Art. 26, Art. 27 FRIA, and the compliance deadlines that become active immediately

08	Your Screening Record	REG-AIMS-CLASS-001
	The completed assessment template — field by field, with explanatory notes and a worked example you can adapt for your own systems	

CHAPTER 01

The Classification *Problem*

Most organisations have answered the classification question by not answering it. The EU AI Act does not accept that.

CRITICAL GAP

ART. 6(2)

UNASSESSED ≠ NOT HIGH-RISK

01

In October 2026, a panel client sends a new section of its AI governance questionnaire: *Has your organisation assessed whether any AI systems used in service delivery constitute High-Risk AI systems under EU AI Act Annex III?*

Two types of organisations receive this question. The first opens its classification register, copies four assessment rationales into the response, and closes section 7(c) in under thirty minutes. The second sends a reply that begins: *"We do not believe our AI tools are high-risk."*

WARNING — UNASSESSED IS NOT NOT HIGH-RISK

An organisation that has not performed an Annex III assessment is not in a "Not High-Risk" position. It is in an **unassessed** position. The EU AI Act does not classify silence. Article 6(2) requires a documented assessment per system. If a system is subsequently found to fall within Annex III, the absence of any prior assessment is an aggravating factor in any regulatory inquiry — not a neutral one.

The peer firm's response — "we don't believe our tools are high-risk" — is not a classification. It is a statement of opinion with no legal basis. The next question from any competent client, regulator, or auditor will be: *On what basis?*

What Makes This Different from Other Compliance Exercises

Annex III classification is the entry gate to the EU AI Act's primary obligation regime. Every downstream obligation — risk management, data governance, technical documentation, human oversight, incident notification, the Fundamental Rights Impact Assessment — flows from a single determination: is this AI system High-Risk under Annex III?

Getting that determination wrong in either direction carries real cost. Wrongly classify a system as High-Risk when it is not, and you impose Arts. 9–15 and Art. 26 obligations on a system that is not subject to them — substantial compliance cost for no legal reason. Wrongly classify a High-Risk system as Not High-Risk, and you leave obligations unmet on a system that is actively regulated — enforcement risk, client trust risk, and the reputational damage of a wrong answer becoming public.

THE RISK OF OVERCLAIMING

False High-Risk

Applying Arts. 9–15 and Art. 26 to systems that are Not High-Risk. Generates disproportionate compliance cost. A common error where Area 8 (administration of justice) is misapplied to solicitor-side legal drafting tools.

THE RISK OF UNDERCLAIMING

False Not High-Risk

Leaving a genuinely High-Risk system unregistered and unmanaged. Triggers enforcement risk at a deadline of 2 December 2027 (standalone AI systems) or 2 August 2028 (embedded in products). The more dangerous error.

The Standard the Act Sets

The EU AI Act does not require you to discover every AI system you use is High-Risk. It requires you to have looked — properly, system by system — and to be able to show your working. The document you are reading is a structured guide to doing that looking correctly, in a way that produces a defensible, signable record.

ITIL 4 — RISK MANAGEMENT AND SERVICE CONFIGURATION MANAGEMENT

Under ITIL 4, the Annex III assessment maps to two practices. As a **Risk Management** input, it determines whether a specific AI system creates a regulatory risk with a defined compliance deadline. As a **Service Configuration Management** record, it documents a property of each Configuration Item (AI system): its EU AI Act classification status. The classification register is the CMDB entry for compliance standing, reviewed annually and on material change — exactly the lifecycle cadence ITIL 4 requires.

CHAPTER 02

Your Legal *Obligation*

Article 6 is the classification gateway. Understanding its three sub-provisions is the prerequisite to every assessment that follows.

ART. 6(1) · (2) · (3)

ANNEX III

ARTS. 9-15 · ART. 26

02

REGULATION (EU) 2024/1689 — ARTICLE 6(2) · THE PRIMARY CLASSIFICATION PROVISION

"In addition to the high-risk AI systems referred to in paragraph 1, AI systems referred to in Annex III shall be considered to be high-risk."

Article 6 contains three operative sub-provisions. For most regulated organisations — law firms, financial institutions, healthcare providers, manufacturers — only one of them is routinely in play.

The Three Sub-Provisions of Article 6

6(1) **The Product Safety Route — Rarely Triggered**

Covers AI systems intended as safety components in products regulated under a list of EU product safety directives — medical devices, machinery, aircraft, elevators, and similar. Triggered when an AI system requires CE marking as part of a regulated product. For organisations using off-the-shelf AI tools (drafting assistants, matter management software, email copilots), this provision is almost never in play.

6(2) **The Annex III Route — The Primary Concern**

The main classification pathway for regulated organisations. Requires assessment of whether the AI system's use case falls within any of the eight Annex III areas. This is the provision that every organisation deploying AI must work through, system by system. It is the subject of this guide.

6(3) **The Narrow Exclusion — Strictly Interpreted**

Provides specific grounds on which a system that would otherwise fall within Annex III may be excluded from High-Risk classification. The exclusion grounds are limited: the system performs only a preparatory task; its output is subject to human override without producing the actual decision; or it detects patterns without making decisions affecting individuals. These are not interchangeable with general human review provisions under Art. 50(4). See Chapter 6 for the full exclusion analysis.

The Obligation Cascade

A High-Risk classification under Art. 6(2) triggers a cascade of obligations. These apply to **deployers** — organisations using High-Risk AI systems in their operations — as well as to providers. The deployer obligations are the ones most regulated organisations will face.

ARTICLE	OBLIGATION	DEPLOYER POSITION	DEADLINE (STANDALONE AI)
Art. 9	Risk Management System	Establish and maintain a documented risk management process for the High-Risk system throughout its lifecycle	2 Dec 2027
Art. 10	Data and Data Governance	Ensure training, validation, and testing data practices meet quality and governance standards	2 Dec 2027
Art. 11	Technical Documentation	Maintain technical documentation describing the system's design, intended purpose, and compliance measures	2 Dec 2027
Art. 12	Record-Keeping	Enable automatic logging of events throughout the system's operational lifetime	2 Dec 2027
Art. 13	Transparency to Users	Ensure sufficient transparency to enable deployers to interpret the system's output and use it appropriately	2 Dec 2027
Art. 14	Human Oversight	Implement human oversight measures enabling effective monitoring, intervention, and override of the system's operations	2 Dec 2027
Art. 15	Accuracy, Robustness, Cybersecurity	Achieve appropriate levels of accuracy and resilience, with cybersecurity measures throughout the lifecycle	2 Dec 2027
Art. 26	Deployer Obligations	Use the system per provider instructions; assign competent human oversight; monitor operation; report incidents to the provider	2 Dec 2027
Art. 27	Fundamental Rights Impact Assessment	Conduct a FRIA before deploying the High-Risk AI system in contexts affecting natural persons	Before deployment

CAUTION – COMPLIANCE DEADLINES ARE ACTIVE

The 2 December 2027 deadline applies to standalone AI systems. AI systems embedded in regulated products that already require third-party conformity assessment have a later deadline of 2 August 2028. If a system is found to be High-Risk, both the classification date and the applicable deadline shall be recorded in the classification register immediately.

CHAPTER 03

Before You *Begin*

Scope your AI inventory correctly. The assessment process is per-system, not per-firm, per-department, or per-product family.

AI INVENTORY

SCOPE DEFINITION

PER-SYSTEM ASSESSMENT

03

What Counts as an AI System for These Purposes

The EU AI Act defines an AI system as a machine-based system designed to operate with varying levels of autonomy that infers, for a given set of objectives, outputs such as predictions, content, recommendations, or decisions. This covers a broad range of tools in routine use across regulated organisations.

TYPICALLY IN SCOPE

Systems to Assess

Generative AI drafting tools (GPT-4o, Claude, Gemini); productivity copilots (M365 Copilot, Google Workspace AI); matter or case management AI; HR screening or performance tools; client intake automation; document review and classification tools; AI-powered risk or credit scoring.

TYPICALLY OUT OF SCOPE

Standard Software

Rule-based automation (macro scripts, if-then workflows); traditional statistical models used solely for internal analytics without individual-affecting outputs; search engines; spam filters; basic spell-check. The boundary can be unclear — when in doubt, assess and document the scope decision.

KEY PRINCIPLE – ONE ASSESSMENT ROW PER SYSTEM

Article 6(2) does not provide a de minimis threshold. A general collective statement — "none of our AI tools are high-risk" — is not an Annex III assessment. Each AI system in use by the organisation requires its own documented assessment row. **Four systems in use means four assessments to complete and retain.** The companion Excel template (REG-AIMS-CLASS-001) provides one row per system and stores the assessment outcome, rationale, and approval.

The Pre-Assessment Checklist

01 Complete your AI system inventory

List every AI system in operational use. Include vendor-supplied tools, embedded AI in software platforms, and any AI developed or procured specifically for the organisation. This is your assessment scope. Do not assess systems that are not in use — but do not exclude systems because you assume they are low-risk.

02 Record the primary use case for each system

The Annex III screening is use-case specific, not system-specific. The same AI model used for two different purposes may receive different classifications. M365 Copilot used for email drafting

and M365 Copilot configured to assist HR performance reviews shall be assessed separately against their respective use cases.

03 **Identify the deployment context**

Note whether each system interacts with individuals outside the organisation (clients, patients, applicants) or operates solely in internal workflows. Systems that generate outputs affecting individuals are more likely to touch Annex III. Internal-only analytics tools are less likely — but still require assessment.

04 **Assign an assessment owner and a sign-off authority**

Each assessment requires a named owner (the person conducting it) and a sign-off authority (the person with authority to approve the classification outcome). In regulated firms, the sign-off authority is typically the COLP, CIO, DPO, or equivalent. The completed record shall be countersigned and dated before it is treated as final.

05 **Set your retention period**

Classification records are compliance records. Retain each assessment record for a minimum of 7 years from the date of sign-off, or for the duration of the AI system's operational life plus 7 years. Update the record when the system's use case changes materially, and trigger a review at least annually.

CHAPTER 04

The Three-Step *Classification Test*

A structured decision process that produces a defensible, signable classification record for any AI system in your inventory.

STEP 1: AREA SCREENING

STEP 2: USE-CASE SPECIFICITY

STEP 3: ART. 6(3) EXCLUSION

04

The classification process is a three-step structured test. The steps are sequential and gated — each step is only reached if the previous step produces a finding that requires further assessment. Each step produces a documented answer. The three answers together constitute the classification record.

CLASSIFICATION DECISION FRAMEWORK – THREE STEPS

Work through each step in sequence. Stop when a step produces a definitive outcome.

STEP 1 – ANNEX III AREA SCREENING (REQUIRED FOR ALL SYSTEMS)

Does the AI system's primary use case match any of the eight Annex III areas?

NO AREA IDENTIFIED → STOP HERE

Document which areas were screened and why none applies to this system's use case. The assessment closes as **Not High-Risk at Step 1**. Record the rationale, obtain sign-off, and set the next annual review date. Step 2 is not required.

AREA IDENTIFIED → PROCEED TO STEP 2

Record which Annex III area was identified as potentially applicable and why. Document the most plausible area — even if you believe the Step 2 result will be Not High-Risk. Proceed to Step 2 to test the use-case specificity.

STEP 2 – USE-CASE SPECIFICITY TEST (ONLY IF STEP 1 IDENTIFIES AN AREA)

Does the organisation's specific use case match one of the listed use cases within the identified Annex III area — not just the general sector?

USE CASE NOT LISTED → NOT HIGH-RISK AT STEP 2

The area is identified but the specific use case does not match any listed use case within that area. Document the area screened, the listed use cases tested, and why the organisation's actual use case does not match. Assessment closes as **Not High-Risk at Step 2**. Obtain sign-off. Step 3 is not required.

USE CASE LISTED → PROCEED TO STEP 3

The organisation's use case matches a listed Annex III use case. Record the specific listed use case and the factual basis for the match. Proceed to Step 3 to assess whether an Art. 6(3) exclusion is available.

STEP 3 – ART. 6(3) EXCLUSION ASSESSMENT (ONLY IF STEPS 1 AND 2 BOTH POINT HIGH-RISK)

Does the system qualify for one of the Art. 6(3) narrow exclusion grounds?

EXCLUSION APPLIES → NOT HIGH-RISK UNDER ART. 6(3)

Record the specific exclusion ground and the factual basis. The assessment closes as **Not High-Risk under Art. 6(3)**. Obtain qualified legal review of the exclusion argument before finalising — Art. 6(3) exclusions are strictly interpreted and the factual basis must be robust.

NO EXCLUSION AVAILABLE → HIGH-RISK CLASSIFICATION

No Art. 6(3) exclusion applies. The system is **High-Risk**. Open Arts. 9–15 obligation rows immediately. Prioritise Art. 27 FRIA before deployment. The 2 December 2027 (standalone) or 2 August 2028 (embedded) compliance deadline is now active. Obtain qualified legal advice before finalising.

NOTE – THE ASSESSMENT IS USE-CASE SPECIFIC, NOT SYSTEM-SPECIFIC

If an AI system is used for multiple distinct purposes within the same organisation, each significant use case shall be assessed separately. A system that is Not High-Risk in one use case may be High-Risk in another. The classification register shall contain one row per system per assessed use case — not one row per system regardless of how it is deployed.

CHAPTER 05

The Eight Annex III Areas — *Area by Area*

Classification logic, explanatory notes, and real-world examples for each area.
Work through the eight areas for every system in your inventory.

6 CLEAR EXCLUSIONS

2 REQUIRE SCRUTINY

AREA 8 MOST MISCLASSIFIED

05

Annex III lists eight areas. For each area, assess whether your AI system's *primary use case* matches one of the specific listed use cases within that area. A general sector match is not sufficient — the specific use case must correspond. Work through all eight areas for every system, even if you expect the outcome to be Not High-Risk.

AREA 1

**Biometrics — Identification and
Categorisation of Natural Persons**

NOT TRIGGERED (STANDARD DEPLOYMENTS)

What it covers: AI systems used for real-time and post-remote biometric identification of natural persons; biometric categorisation systems inferring or deducing sensitive attributes (race, political opinions, trade union membership, religious beliefs, sexual orientation); and emotion recognition systems.

Classification logic: This area is not triggered by general-purpose AI tools such as drafting assistants, matter management software, email copilots, or document analysis tools. None of these tools perform biometric identification or emotion recognition. Assessment closes at Step 1 — **Not High-Risk**.

When to look closer: Organisations using AI-powered access control systems, video surveillance with automated identification, or any system that infers personal characteristics from physiological data should give Area 1 close attention. Facial recognition, voice biometric authentication, and gait analysis tools all fall within Area 1 scope.

EXAMPLE CLASSIFICATION — WORKED

Generative AI Drafting Tool (e.g. GPT-4o)

Area 1 screened. Use case: legal document drafting and research summarisation. The system does not perform biometric identification, categorisation of sensitive attributes, or emotion recognition. Area 1: not applicable. Assessment closes — **Not High-Risk at Step 1**.

AREA 2

Critical Infrastructure — Management and Operation

NOT TRIGGERED (STANDARD DEPLOYMENTS)

What it covers: AI systems intended to be used as safety components in digital infrastructure, road traffic, water supply, gas, heating, and electricity supply. The key qualifier is "safety component" — the AI must be integral to the safe operation of the infrastructure, not merely used by an organisation that is part of a supply chain.

Classification logic: An organisation using AI to draft documents, manage clients, or analyse data is not managing or operating critical infrastructure. The connection must be direct and operational. Assessment closes at Step 1 for all standard knowledge-economy AI tools — **Not High-Risk**.

When to look closer: Energy companies, water utilities, and telecoms providers using AI to manage operational safety systems — load balancing, pressure monitoring, network fault detection — should assess Area 2 carefully. The test is whether the AI system's failure or malfunction would create a safety risk to the infrastructure itself.

AREA 3

Education and Vocational Training

NOT TRIGGERED (STANDARD DEPLOYMENTS)

What it covers: AI systems used to determine access to or assignment within educational institutions; AI used to evaluate the learning outcomes of students and assess their competences; AI used to detect prohibited behaviour in assessment contexts (e.g., exam monitoring).

Classification logic: Regulated businesses are not educational institutions for this purpose. Internal training programmes that use AI to track completion or test knowledge do not constitute educational assessment systems under Area 3 — the use case must be one where the AI system's output directly determines access to or standing within an educational institution. Assessment closes at Step 1 for standard business AI deployments — **Not High-Risk**.

When to look closer: Universities, colleges, vocational training providers, and online education platforms using AI to assess student performance, allocate places, or monitor examinations shall give Area 3 specific attention. The scope is narrower than "any AI used in education" — it targets AI that determines access and outcomes.

AREA 4

Employment, Worker Management, and Access to Self-Employment

REVIEW USE CASE

What it covers: AI systems used for recruitment or selection — particularly CV sifting and interview screening; AI used to make decisions affecting working conditions, task allocation, and performance monitoring; AI used to assess, promote, or terminate employment relationships; AI used in access to self-employment or contracting.

Classification logic: Area 4 is the one area where most regulated organisations' internal AI deployments require closer scrutiny. Productivity copilots used for document drafting are not employment management systems. However, the same copilot configured to assist HR in producing performance ratings or disciplinary reports that directly influence employment decisions moves into Area 4 territory.

The test: Does the AI system's output directly inform a decision that affects an individual's employment status, working conditions, or access to work? If yes, the use-case specificity test in Step 2 applies and the use case shall be compared against the listed Area 4 use cases in detail.

CAUTION – HR AI TOOLS REQUIRE CAREFUL SCOPING

Many HR platforms now include AI features as standard. Automated CV filtering, AI-assisted interview scoring, workforce analytics generating individual performance flags, and task allocation algorithms that affect shift patterns or workloads shall all be assessed against Area 4. The assessment shall be documented for each distinct HR AI use case, not for the HR platform as a whole.

EXAMPLE CLASSIFICATION — CONDITIONAL NOT HIGH-RISK

M365 Copilot — Document Drafting (Area 4)

Area 4 screened. Use case: draft assistance for legal documents and correspondence in matter delivery. The system does not produce or directly inform employment decisions — its output is legal text reviewed and deployed by solicitors. Area 4: not applicable to this use case. Assessment closes — **Not High-Risk at Step 1**. *Conditional caveat:* if the same tool is configured to assist HR in drafting performance assessments or disciplinary records, the Area 4 assessment shall be reopened for that use case.

AREA 5

Access to Essential Private Services and Public Services and Benefits

REVIEW USE CASE

What it covers: AI systems used to evaluate the creditworthiness of natural persons or to establish their credit scores; AI used to determine access to or the terms of public benefit schemes; AI used to assess risk in health insurance, life insurance, and similar products; AI used in emergency services dispatching.

Classification logic: This area is most commonly triggered for financial services firms, insurers, and public sector bodies. The core question is whether the AI system produces an assessment that directly determines a natural person's access to a financial product, benefit, or essential service — not merely whether the organisation operates in a regulated sector.

The test: Does the AI system generate a creditworthiness score, an insurance risk assessment, or an eligibility recommendation that affects whether a natural person receives a service or benefit? Matter management software that tracks billing and deadlines does not trigger Area 5. An AI tool that generates credit risk flags used in lending decisions does.

EXAMPLE CLASSIFICATION – CONDITIONAL NOT HIGH-RISK

LEAP AI — Matter Management (Area 5)

Area 5 screened. Use case: matter management, time recording, billing, and file workflow. The system does not produce creditworthiness assessments, determine client eligibility for services, or generate insurance risk ratings. Area 5: not applicable to this use case. Assessment closes — **Not High-Risk at Step 1**. *Conditional caveat:* if LEAP AI is configured to make autonomous recommendations on matter acceptance, fee arrangements, or client eligibility that are acted upon without human review, the Area 5 assessment shall be reopened.

AREA 6

Law Enforcement

NOT TRIGGERED (STANDARD DEPLOYMENTS)

What it covers: AI systems used by law enforcement authorities — police, prosecutors, intelligence services — to assess risk posed by individuals, detect emotions in criminal investigations, analyse CCTV or surveillance data, predict criminal activity, detect fraud in criminal contexts, and profile individuals in crime-related investigations.

Classification logic: The key qualifier is "competent law enforcement authorities." A law firm, financial institution, or healthcare provider is not a law enforcement authority. Acting as legal representative for police forces, defendants, or prosecuting authorities does not bring the organisation within Area 6. Use of AI to detect internal fraud or compliance breaches does not trigger Area 6 — that use case falls under the organisation's own risk management, not law enforcement. Assessment closes at Step 1 — **Not High-Risk** for all standard regulated business AI deployments.

AREA 7

Migration, Asylum, and Border Control Management

IMMIGRATION ORGANISATIONS: REVIEW

What it covers: AI systems intended to be used by competent public authorities to assess risk, detect forgeries, assist in examination of applications, and make decisions in migration, asylum, and border control contexts. Includes AI used to verify travel documents and assess the reliability of information provided by applicants.

Classification logic: Again, the key qualifier is "competent public authorities." A law firm practising immigration law, or an employer managing visa sponsorships, is not a competent public authority. The organisation's use of AI in preparing immigration applications or advising clients on immigration status does not trigger Area 7 — the provision covers the authority making the decision, not the professional advising the applicant.

CAUTION — IMMIGRATION FIRMS AND BORDERLINE CASES

If an AI tool is used to assist in the preparation of asylum applications and its output is of a nature that directly influences how a competent authority determines status — for example, automated generation of country condition assessments or risk profiling of individual circumstances — this warrants specific legal analysis before concluding Not High-Risk. Borderline cases shall be recorded as "Requires Legal Advice" in the classification register and referred for external review before finalising the assessment.

AREA 8

Administration of Justice and Democratic Processes

MOST MISCLASSIFIED AREA

What it covers: AI systems intended to be used by a judicial or quasi-judicial body to assist in researching and interpreting facts and the law, and in applying the law to a concrete set of facts. Also covers AI intended to influence electoral and voting processes.

The critical distinction — adjudication vs. advocacy: The phrase "judicial or quasi-judicial body" is the operative qualifier. A law firm, in-house legal team, or legal technology provider is *not* a judicial or quasi-judicial body. The provision covers AI systems used by courts, tribunals, arbitration panels, and similar bodies *in the exercise of their adjudicative function*. It does not cover AI used by solicitors, barristers, or legal teams to prepare material for those bodies.

A solicitor using GPT-4o to research case law, draft a skeleton argument, summarise a witness statement, or review contractual clauses is engaged in legal advocacy — preparing material for submission to a judicial body. The court's own use of AI in its processes may fall within Area 8. The solicitor's preparation of material for that court does not.

WARNING — AREA 8 MISCLASSIFICATION IS COMMON AND COSTLY IN BOTH DIRECTIONS

A firm that classifies its legal drafting AI as High-Risk under Area 8 on the basis that it "relates to" the administration of justice has misread the Annex. The obligations that flow from that misclassification — Arts. 9–15, Art. 26, Art. 27 FRIA — are substantial and disproportionate to a genuinely Not High-Risk tool. The correct classification for solicitor-side legal drafting, research, and document review tools is **Not High-Risk**, with a documented rationale that explains precisely why Area 8 does not apply — the firm is not a judicial body, and the AI system is not assisting in adjudication.

EXAMPLE CLASSIFICATION — AREA 8 CORRECTLY APPLIED

GPT-4o — Legal Research and Drafting (Area 8)

Area 8 screened. Use case: researching case law, drafting skeleton arguments, summarising evidence, and producing legal correspondence. The system is used by solicitors — not by a judicial or quasi-judicial body — and its outputs are prepared for submission to courts, not used in the exercise of judicial functions. Area 8 listed use case test: the Annex covers AI assisting bodies in applying law to facts in an adjudicative capacity. This use case is legal advocacy, not adjudication. Area 8: not applicable. Assessment closes — **Not High-Risk at Step 2.**

ITIL 4 — COMPLETING THE AREA-BY-AREA SCREEN

Under ITIL 4 Configuration Management, the area-by-area screen constitutes the **Identification** activity for each Configuration Item (AI system): establishing what properties the item has, what its compliance characteristics are, and what its regulatory status is. This screen is not a one-time event — it is a lifecycle record that is updated on material change and reviewed annually. Each area screened, and the outcome of that screen, is a field in the system's configuration record.

CHAPTER 06

Article 6(3) — The *Narrow Exclusion*

Only reached if Steps 1 and 2 both point toward High-Risk. The exclusion is strictly interpreted. Understanding what it is — and what it is not — prevents the most common classification error.

ART. 6(3)

NOT A GENERAL ESCAPE HATCH

DISTINCT FROM ART. 50(4)

06

REGULATION (EU) 2024/1689 — ARTICLE 6(3) · EXCLUSION GROUNDS

"Notwithstanding paragraph 2, an AI system referred to in Annex III shall not be considered high-risk if it does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision-making."

Article 6(3) is only reached if Steps 1 and 2 both produce affirmative results — the system falls within an Annex III area, and the specific use case matches a listed use case within that area. At that point, the exclusion grounds in Art. 6(3) are assessed.

The Three Documented Exclusion Grounds

6(3) · a **Preparatory Task Only**

The AI system is intended to perform a narrow preparatory task that is not, by itself, used to make or substantially contribute to a decision that directly affects natural persons. The system's output is an input to a human decision-making process, and the human decision-maker has full discretion. This ground requires that the system's role be genuinely preparatory — not that a human is nominally involved in reviewing output that the organisation in practice treats as determinative.

6(3) · b **Human Override of Output**

The AI system's output is subject to human override before it produces any decision, and the human override is genuine and practically available — not a formality. The system must be designed so that the human can meaningfully review, modify, and reject the AI output. The override must be documented and habitual, not merely theoretically possible.

6(3) · c **Pattern Detection Without Individual-Affecting Decisions**

The AI system is used only to detect patterns, anomalies, or trends in data, and its outputs do not directly determine decisions that affect individual natural persons. A fraud detection system that flags anomalous transactions for human review — without itself making the decision to freeze an account or reject a transaction — may qualify under this ground, provided the human review is genuine and the system's flag is not treated as determinative in practice.

WARNING — "WE HAVE HUMAN REVIEW" IS NOT AN ART. 6(3) EXCLUSION

The most common classification error at Step 3 is asserting Art. 6(3) on the basis of general human review of AI outputs. This conflates two distinct legal provisions. Art. 50(4) addresses the transparency labelling obligation for AI-generated content — it is satisfied by disclosing that content is AI-generated and by having editorial procedures. Art. 6(3) addresses classification — it requires that the system's output does not materially influence the outcome of a decision affecting natural persons. A firm that asserts Art. 6(3) because "our solicitors review all AI output" has not engaged with the legal standard. The question is not whether a human sees the output — it is whether the output, in practice, materially determines the decision.

Documenting an Art. 6(3) Claim

If an Art. 6(3) exclusion is asserted, the following shall be documented in the classification record:

i The specific exclusion ground relied upon

Identify which of the three grounds applies — preparatory task only, genuine human override, or pattern detection without individual-affecting decisions. A general assertion of "not high-risk" is insufficient. The specific ground shall be named.

ii The factual basis for the exclusion

Describe in specific terms how the system's design and deployment satisfies the identified exclusion ground. Where the human override ground is relied upon, evidence of the override mechanism and its practical use shall be noted. Where the preparatory task ground is relied upon, the scope of the preparatory function and the scope of the human decision it feeds shall be described.

iii Legal review

An Art. 6(3) exclusion assertion shall be reviewed by qualified legal counsel before the classification record is finalised. The classification produces a legal conclusion — that a system which would otherwise be High-Risk is excluded from that category. That conclusion shall not rest solely on internal assessment without legal sign-off.

CHAPTER 07

If You Find a *High-Risk System*

The obligation cascade is immediate. The compliance deadlines are fixed. Here is what to do in the first 30 days after a High-Risk classification.

HIGH-RISK CONFIRMED

ARTS. 9-15 ACTIVE

2 DEC 2027 / 2 AUG 2028

07

A High-Risk classification is a compliance trigger, not a compliance failure. The Act does not prohibit deploying High-Risk AI systems — it requires that the deployment be managed to defined standards by defined deadlines. The response to a High-Risk finding is structured, not reactive.

First 30 Days — Immediate Actions

01 Record the classification formally and obtain legal sign-off

Enter the High-Risk classification in the classification register. Record the area, the specific listed use case, the Step 3 outcome, the classification rationale, and the date. Obtain legal counsel review of the classification before treating it as final — a High-Risk determination has significant compliance consequences and shall be defensible.

02 **Open Arts. 9–15 obligation rows in your compliance tracker**

Each Article 9–15 obligation becomes a tracked compliance item for this specific system. The companion REG-AIMS-CLASS-001 template generates one row per obligation per High-Risk system automatically. Each row shall have a named owner, a target completion date within the 2 December 2027 or 2 August 2028 deadline, and a status field updated at least quarterly.

03 **Trigger the Fundamental Rights Impact Assessment (Art. 27 FRIA)**

The FRIA is required before deploying a High-Risk AI system in contexts that affect natural persons. If the system is already deployed, the FRIA is required as a matter of urgency — the FRIA Workbook (EBOOK-AIMS-FRIA-001) provides the template and guidance. The FRIA cannot be deferred until the Art. 9–15 compliance programme is complete — it is a prerequisite to lawful deployment, not a downstream deliverable.

04 **Review the Art. 26 deployer obligations**

Article 26 requires deployers to use High-Risk AI systems in accordance with the provider's instructions; assign named human oversight to natural persons with the competence to monitor the system; implement procedures to report incidents and serious malfunctions to the provider; and inform relevant staff of the system's High-Risk classification and its implications.

05 **Consider the regulatory sandbox pathway**

If the High-Risk classification is borderline — or if the organisation is uncertain how the Art. 9–15 obligations should be implemented for a specific system — the national AI regulatory sandbox (available from 2 August 2027 under Omnibus VII) provides a supervised testing environment where competent authorities can provide guidance. Note the sandbox pathway in the obligation tracker for borderline systems.

NOTE – REGULATORY SANDBOXES FROM 2 AUGUST 2027

Omnibus VII (adopted 29 June 2026) postponed national AI regulatory sandboxes to **2 August 2027**. Sandboxes allow supervised testing of AI systems under competent authority oversight — directly relevant where a High-Risk classification produces implementation questions that cannot be resolved internally. The sandbox pathway is a risk-managed route to regulatory clarity before the 2 December 2027 compliance deadline. Note this option in the obligation record for any system where the implementation path is uncertain.

CHAPTER 08

Your *Screening Record*

The completed assessment template — field by field, with explanatory notes and a worked example you can adapt for every system in your inventory.

REG-AIMS-CLASS-001

COMPANION EXCEL TEMPLATE

7-YEAR RETENTION

08

The screening record is the document that converts the three-step assessment into a signed, dated, retainable compliance record. It is the answer to the client questionnaire, the regulatory inquiry, and the internal audit. Complete one record per system per assessment event. The companion REG-AIMS-CLASS-001 Excel template provides the structure — this chapter explains what goes in each field.

Section A — Document Control

REG-AIMS-CLASS-001 — DOCUMENT CONTROL BLOCK	TEMPLATE
Document ID	
REG-AIMS-CLASS-001 (fixed — do not change)	
Organisation Name	
[Your organisation's full legal name]	
Registration Number	
[Company number / SRA number / FCA reference as applicable]	
Version	
1.0 (increment on each material update)	
Issue Date	
[Date of first issue — not the date of this assessment row]	
Next Review Date	
[12 months from issue date, or earlier if material change]	
Document Status	
ACTIVE / UNDER REVIEW / SUPERSEDED	
Retention Period	
7 years from issue date, or system lifetime + 7 years	

Section B — AI System Identification

PER-SYSTEM ASSESSMENT — SYSTEM IDENTIFICATION	ONE ROW PER SYSTEM
System Reference	
AIMS-SYS-001 (from your AI system inventory)	
System Name	
[Common name of the AI tool, e.g. GPT-4o, M365 Copilot Word]	
Provider	
[Name of the provider/vendor, e.g. OpenAI, Microsoft]	
Primary Use Case	
[Specific description of how the system is used in your organisation — not the vendor's description of the product's capabilities]	
Deployment Context	
[Internal-only / Client-facing / Hybrid. Note whether outputs are used in individual-affecting decisions]	
Assessment Owner	
[Name and role of the person conducting this assessment]	

Section C — The Three-Step Assessment

STEP 1 – ANNEX III AREA SCREENING

Areas Screened

List all eight areas: Area 1 (Biometrics), Area 2 (Critical Infrastructure), Area 3 (Education), Area 4 (Employment), Area 5 (Essential Services), Area 6 (Law Enforcement), Area 7 (Migration), Area 8 (Administration of Justice). Record each area considered.

Area Identified

[Name of area identified as potentially applicable, or "None" if no area applies]

Step 1 Rationale

[Plain English explanation of why the identified area was selected, or why no area applies. Two to four sentences minimum.]

Step 1 Result

No area identified — assessment closes here (Not High-Risk at Step 1) / Area identified — proceed to Step 2

STEP 2 – USE-CASE SPECIFICITY TEST (COMPLETE ONLY IF STEP 1 IDENTIFIED AN AREA)

Use Case Tested

[The specific listed use case within the identified Annex III area that was tested against the organisation's actual use case]

Organisation's Use Case

[Restate the organisation's actual use case as assessed against the listed use case above]

Step 2 Rationale

[Plain English explanation of whether and why the organisation's use case matches or does not match the listed Annex III use case. Three to six sentences minimum.]

Step 2 Result

Use case does not match — assessment closes here (Not High-Risk at Step 2) / Use case matches — proceed to Step 3

STEP 3 — ART. 6(3) EXCLUSION ASSESSMENT (COMPLETE ONLY IF STEPS 1 AND 2 BOTH POINT HIGH-RISK)
Exclusion Ground
<i>Preparatory task only / Genuine human override / Pattern detection without individual-affecting decisions / No exclusion available</i>
Exclusion Factual Basis
<i>[Where an exclusion is asserted: the specific factual basis — how the system's design and deployment satisfies the identified exclusion ground. Where no exclusion applies: state that no ground is available and why.]</i>
Legal Review
<i>[Name of legal counsel who reviewed the Art. 6(3) assessment, and date of review — required where exclusion is asserted]</i>
Step 3 Result
<i>Exclusion applies — Not High-Risk under Art. 6(3) / No exclusion available — High-Risk classification confirmed</i>

Section D — Classification Outcome and Sign-Off

CLASSIFICATION OUTCOME — TO BE COMPLETED FOR EVERY ASSESSMENT
Classification
NOT HIGH-RISK (at Step []) / HIGH-RISK (Annex III — Area [])
Classification Rationale
[A plain-English summary of the assessment — one paragraph, three to six sentences — suitable for quoting verbatim in a client questionnaire or regulatory inquiry. This is the answer to "On what basis?"]
Conditional Caveats
[Any conditions under which this Not High-Risk conclusion would require reassessment — e.g. "if configured to assist HR performance decisions" or "if deployed in client-facing credit assessment"]
Assessment Date
[Date this assessment was completed]
Approved By
[Name and role of sign-off authority — COLP / DPO / CIO / equivalent]
Approval Date
[Date of approval — may differ from assessment date]
Next Review Date
[12 months from approval date, or earlier if material change to use case or deployment]

THE COMPANION EXCEL TEMPLATE

REG-AIMS-CLASS-001 is available as a companion Excel download with this guide. It provides the complete five-table structure: firm identification, per-system registry, assessment record, High-Risk obligation tracker, and audit log. The template generates one assessment row per system and includes conditional formatting that flags overdue reviews, missing sign-offs, and High-Risk systems with outstanding obligation rows. The Excel template is the working document — this chapter is the explanatory guide to completing it correctly.

Worked Example — Complete Assessment Record

WORKED EXAMPLE — GPT-4o AT A LEGAL SERVICES FIRM

AIMS-SYS-001 · GPT-4o · Legal Research and Drafting

Primary use case: Generative AI drafting and research tool used by solicitors to research case law, draft legal correspondence and skeleton arguments, summarise witness statements, and review contractual clauses. Output is reviewed by qualified solicitors before use.

Step 1: Eight Annex III areas screened. Area 8 (Administration of Justice) identified as the most plausible area on the basis that the tool is used in connection with legal proceedings. Area identified. Proceed to Step 2.

Step 2: Area 8 listed use case tested: AI systems intended to be used by a judicial or quasi-judicial body to assist in researching and interpreting facts and the law and in applying the law to a concrete set of facts. Organisation's use case: AI used by solicitors (a law firm) to prepare material for submission to courts. The firm is not a judicial or quasi-judicial body; it is an advocate before such bodies. The Act covers adjudication, not advocacy. Use case does not match the listed use case. Assessment closes — **Not High-Risk at Step 2.**

Classification rationale: GPT-4o is used at this firm by solicitors engaged in legal advocacy — preparing documents for submission to courts and tribunals. Annex III Area 8 covers AI systems assisting judicial and quasi-judicial bodies in the exercise of their adjudicative function. This firm is not a judicial body and does not exercise adjudicative functions. The AI system's use is preparatory advocacy, not adjudication. Area 8 does not apply. All other seven Annex III areas screened and found not applicable to this use case. Classification: Not High-Risk.

Conditional caveat: If the system is deployed within a judicial body, quasi-judicial body, or arbitration panel in an adjudicative capacity, this classification shall be reopened.

Approved by: COLP · July 2026 · Next review: July 2027

WHAT COMES NEXT

The Assessment Is *Done*. The Record Is the Asset.

The classification record you have just completed is not a filing exercise. It is the answer to the question every panel client, regulator, and auditor will ask: *how do you know your AI systems are not high-risk?* The difference between a firm that has done this work and a firm that has not becomes visible the moment that question arrives.

UNUS LONDON · GOVERNANCE ACADEMY

Take the Next Step

If your classification exercise identifies one or more High-Risk systems, the Fundamental Rights Impact Assessment is your next obligation. The FRIA Workbook walks you through the statutory process, field by field, with explanatory commentary alongside every section.

unuslondon.com/governance-academy

DOCUMENT ID: EBOOK-AIMS-CLASS-001 V1.0 · SEPTEMBER 2026
LEGAL BASIS: REGULATION (EU) 2024/1689 · ART. 6(1)(2)(3) · ANNEX III
COMPANION: REG-AIMS-CLASS-001 ANNEX III CLASSIFICATION REGISTER (EXCEL)
RELATED: EBOOK-AIMS-FRIA-001 – THE EU AI ACT FRIA WORKBOOK
STANDARDS: ISO/IEC 42001:2023 · ITIL 4 · IEC 82079-1:2012
© UNUS LONDON LTD · ALL RIGHTS RESERVED · NOT LEGAL ADVICE