EUROPEAN COMPLIANCE SUITE

EU AI Act Compliance Evidence Templates

EXPLANATORY NOTICE AND TEMPLATE

Fundamental Rights Impact Assessment for high-risk AI systems, required by Article 27 of Regulation (EU) 2024/1689 (AI Act)

Template reference	ECS-AIA-HR-FRIA
Version	1.0
Issued	July 2026
Obligation addressed	Article 27 AI Act — assessment of the impact on fundamental rights that use of a high-risk AI system may produce, and notification of its results
Addressee	Deployers that are bodies governed by public law or private entities providing public services, and deployers of Annex III point 5(b) and 5(c) systems
Position in the pack	The flagship assessment. Performed prior to first use, after ECS-AIA-HR-SCOPE has confirmed the trigger
Companion templates	ECS-AIA-HR-SCOPE — applicability and trigger; ECS-AIA-HR-RECORD — Article 26 obligations record; ECS-AIA-HR-OVERSIGHT — human oversight, logging and monitoring evidence

This Explanatory Notice and the Template annexed to it are issued by European Compliance Suite as private compliance documentation. They are not a Union instrument, are not endorsed by the European Commission or the AI Office, and do not constitute legal advice. Where the Commission adopts an official template, guidance or implementing act covering the same obligation, that instrument prevails.

1. Background and purpose

(1) Article 27 of Regulation (EU) 2024/1689 (the 'AI Act') requires certain deployers of high-risk AI systems, prior to deployment, to perform an assessment of the impact on fundamental rights that the use of the system may produce. The provision sets out six elements the assessment must contain, requires the deployer to notify the market surveillance authority of the results, and provides that the assessment applies to the first use of the system and must be updated where its elements cease to be current.

(2) The assessment is unusual among AI Act obligations in that its subject-matter is not the system but the deployment. The same system used by two organisations, in two processes, affecting two populations, produces two different assessments. Nothing the provider supplies can substitute for it, because the provider does not know the deployer's process, the population it serves, or the consequences for that population of an adverse outcome. What the provider supplies — in particular the information required under Article 13 — is an input to the assessment, not the assessment itself.

(3) This distinguishes the fundamental rights impact assessment from most compliance documents, and it is why a generic completed example is of limited use. What can be supplied is a structure that follows the statutory elements, prompts that surface the questions organisations tend to avoid, and a discipline for recording reasoning so that the assessment can be defended, updated, and notified.

(4) The assessment is not a formality. Its purpose is to identify, before deployment, the specific risks of harm to identified groups, and to decide what will be done if those risks materialise. An assessment that identifies no risks, or that identifies risks without stating what will be done about them, has not performed the exercise the provision requires.

RELEVANT LEGAL TEXT

Article 27(1) AI Act.

Prior to deploying a high-risk AI system referred to in Article 6(2), with the exception of high-risk AI systems intended to be used in the area listed in point 2 of Annex III, deployers that are bodies governed by public law, or are private entities providing public services, and deployers of high-risk AI systems referred to in points 5 (b) and (c) of Annex III, shall perform an assessment of the impact on fundamental rights that the use of such system may produce. For that purpose, deployers shall perform an assessment consisting of: (a) a description of the deployer's processes in which the high-risk AI system will be used in line with its intended purpose; (b) a description of the period of time within which, and the frequency with which, each high-risk AI system is intended to be used; (c) the categories of natural persons and groups likely to be affected by its use in the specific context; (d) the specific risks of harm likely to have an impact on the categories of natural persons or groups of persons identified pursuant to point (c) of this paragraph, taking into account the information given by the provider pursuant to Article 13; (e) a description of the implementation of human oversight measures, according to the instructions for use; (f) the measures to be taken in the case of the materialisation of those risks, including the arrangements for internal governance and complaint mechanisms.

Article 27(2) to (4) AI Act (extract).

The obligation laid down in paragraph 1 applies to the first use of the high-risk AI system. The deployer may, in similar cases, rely on previously conducted fundamental rights impact assessments or existing assessments carried out by provider. If, during the use of the high-risk AI system, the deployer considers that any of the elements listed in paragraph 1 has changed or is no longer up to date, the deployer shall take the necessary steps to update the information. Once the assessment referred to in paragraph 1 of this Article has been performed, the deployer shall notify the market surveillance authority of its results … If any of the obligations laid down in this Article is already met through the data protection impact assessment conducted pursuant to Article 35 of Regulation (EU) 2016/679 … the fundamental rights impact assessment referred to in paragraph 1 of this Article shall complement that data protection impact assessment.

2. Which fundamental rights are in issue

(5) Article 27 does not enumerate the rights to be assessed. The reference point is the Charter of Fundamental Rights of the European Union, and the rights actually engaged depend on the deployment.

The following are those most commonly in issue in high-risk AI deployments, offered as a prompt rather than a closed list.

Right	Typical relevance in a high-risk deployment
Human dignity (Art 1)	Treatment of the person as an object of scoring or classification; automated handling of matters affecting basic welfare.
Respect for private and family life (Art 7)	Observation, profiling and inference beyond what the person expects or has agreed to.
Protection of personal data (Art 8)	Lawfulness, accuracy, minimisation and the exercise of data subject rights. Overlaps with the data protection impact assessment.
Freedom of expression and information (Art 11)	Chilling effects where the system observes or moderates expression.
Non-discrimination (Art 21)	Differential accuracy or outcomes across protected characteristics, including through proxies. Engaged in almost every deployment affecting individuals.
Rights of the child (Art 24)	Where the affected population includes minors, or decisions affect them indirectly.
Rights of the elderly and of persons with disabilities (Arts 25, 26)	Accuracy and accessibility for populations under-represented in training data or in interface design.
Fair and just working conditions (Art 31)	Employment, task allocation, monitoring and performance evaluation deployments.
Consumer protection (Art 38)	Pricing, eligibility and service allocation decisions.
Right to an effective remedy and to a fair trial (Art 47)	Whether an adverse outcome can be understood, challenged and corrected. Frequently the decisive right in practice.
Good administration (Art 41)	For public bodies: reasons for decisions, the right to be heard, and access to the file.

(6) Two rights deserve particular attention because they are the ones most often assessed superficially. Non-discrimination is engaged wherever a system produces different outcomes for different groups, and the absence of a protected characteristic in the input data does not answer the question, since proxies are common and often unrecognised. The right to an effective remedy is engaged wherever an adverse outcome is difficult to understand or contest; an assessment that describes an appeal route without asking whether the affected person could realistically use it has not addressed the right.

3. The six statutory elements

(7) The Template follows Article 27(1)(a) to (f) in order. Each element is treated below with the questions it requires the deployer to answer, and the failure modes it is designed to prevent.

3.1. The deployer's processes (a)

(8) The description must cover the processes in which the system will be used, in line with its intended purpose. Two failures recur: describing the system rather than the process, and describing the process at a level of generality that conceals where the system's output actually bears on a person. The useful description states what happens before the system is invoked, what it produces, who receives that output, what they do with it, and what the person on the receiving end experiences.

(9) The relationship between the system's output and the eventual decision should be stated precisely. 'Decision support' covers arrangements ranging from a system that ranks cases for human review to

one whose output is followed in substantially all cases. The distinction matters for every subsequent element, and where the practical effect is that the output is rarely departed from, that should be recorded rather than described as advisory.

3.2. Period and frequency of use (b)

(10) The description must cover the period within which and the frequency with which the system is intended to be used. This element establishes scale, and scale bears on risk: a system applied to several hundred cases a year presents a different risk profile from one applied to several million. Pilots and phased rollouts should be described as such, with the intended progression.

3.3. Affected categories and groups (c)

(11) The assessment must identify the categories of natural persons and groups likely to be affected by the use in the specific context. Two distinctions matter. Persons affected are not only the direct subjects of the system's output: family members, dependants, and those affected by a resource allocation decision may be affected without ever being processed. And within the subject population, sub-groups likely to experience different outcomes should be identified expressly, because element (d) requires risks to be assessed against the categories identified here, and a population described only as 'applicants' cannot support a meaningful analysis of differential harm.

(12) Groups warranting specific consideration include those under-represented in the data on which the system was developed, those for whom the system's inputs are less reliable, those with limited capacity to detect or contest an adverse outcome, and those for whom an adverse outcome has consequences of a different order — loss of housing, benefits, liberty or care.

3.4. Specific risks of harm (d)

(13) This is the substantive core of the assessment. The risks must be specific, must be linked to the categories identified under (c), and must take into account the information given by the provider under Article 13 — in particular the system's known limitations, its accuracy and its performance characteristics.

(14) Three disciplines make this element useful. Risks should be expressed as harms to people rather than as failures of the system: not 'the model may be inaccurate' but 'an applicant may be refused and lose access to housing on the basis of an incorrect score'. Risks should be differentiated across the identified groups, since a uniform risk statement usually means the analysis has not been done. And the provider's stated limitations should be confronted directly: where the provider discloses reduced accuracy for a population that forms part of the deployment's affected group, that disclosure creates a risk the deployer knows about and must address.

(15) Risks arising from the system operating as intended should be assessed alongside risks arising from error. A system that performs exactly as designed may still produce harm through the effect of its design on a group, and that possibility is within the scope of the assessment.

3.5. Human oversight measures (e)

(16) The assessment must describe the implementation of human oversight measures, according to the instructions for use. The description should state who exercises oversight, what they are able to see, what they are empowered to do, and what would cause them to intervene. Oversight that consists of a person able to approve or reject an output without access to the reasons for it, or without practical capacity to depart from it, should be recorded as such, because describing it as effective oversight when it is not is the most consequential misstatement an assessment can contain.

(17) The detailed evidence of oversight arrangements is recorded in ECS-AIA-HR-OVERSIGHT; this element records the arrangements as they bear on the risks identified under (d), and should state expressly which risks oversight is expected to mitigate and which it cannot.

3.6. Measures on materialisation, governance and complaints (f)

(18) The assessment must state the measures to be taken if the risks materialise, including the arrangements for internal governance and complaint mechanisms. This element converts the analysis into commitments, and it is the element on which an authority reviewing a notified assessment is most likely to test the organisation.

(19) Useful measures are specific, owned and triggered: they identify what happens, who does it, and what condition sets it in motion. The complaint mechanism should be assessed from the position of the affected person, addressing how they would learn that the system was involved, how they would raise a complaint, what happens to the decision while the complaint is considered, and who resolves it. Where the deployment affects persons who may face difficulty using a standard complaints route, that should be addressed.

4. Relationship with the data protection impact assessment

(20) Where obligations under Article 27 are already met through a data protection impact assessment conducted under Article 35 GDPR, the fundamental rights impact assessment complements that assessment. The two overlap substantially but are not co-extensive: the data protection impact assessment is concerned with risks to rights and freedoms arising from the processing of personal data, while the fundamental rights impact assessment is concerned with the impact of the use of the system on fundamental rights generally, including where no personal data is processed or where the harm arises from the decision rather than the processing.

(21) The practical approach is to conduct them together, to record which document addresses which element, and to cross-refer rather than duplicate. Section 8 of the Template records that mapping. Where the deployer relies on an existing data protection impact assessment for an element of Article 27, the reliance should identify the specific passage relied upon.

5. Notification, timing and updating

(22) The assessment must be performed prior to deploying the system, and applies to its first use. Once performed, the deployer must notify the market surveillance authority of its results, submitting the filled-out template referred to in Article 27(5). The AI Office is to develop a template questionnaire, including through an automated tool, to facilitate compliance in a simplified manner; where that instrument is available at the time of the assessment, the deployer should complete and submit it, and this document then serves as the underlying analysis supporting the notified results.

(23) A deployer may, in similar cases, rely on a previously conducted fundamental rights impact assessment or on an existing assessment carried out by the provider. Reliance should be recorded together with the basis for considering the cases similar; a different affected population, a different process, or a materially different scale will ordinarily defeat similarity.

(24) Where any element ceases to be current during use, the deployer must take the necessary steps to update the information. The Template therefore requires a review cadence and a register of trigger events, and treats the assessment as a maintained document rather than a completed one.

6. Who should perform the assessment

(25) The assessment requires knowledge of the process, of the system’s technical characteristics and limitations, and of the affected population. It is rarely well performed by one function alone. Practice that produces defensible assessments typically involves the process owner, a technical function able to interpret the provider’s Article 13 information, a legal or compliance function, the data protection officer where one is appointed, and — where the deployment affects service users or workers — a route by which their perspective is represented.

(26) The Template records who contributed and in what capacity. Where the perspective of the affected population was not sought, that should be recorded as a limitation rather than omitted, since it bears on the weight the assessment can carry.

7. Retention and enforcement

(27) The assessment and the material supporting it should be retained for at least ten years after the deployment ends, and be capable of production to a national competent authority in a Union language. Superseded versions are retained, since the relevant question on enquiry is what was assessed before the deployment in issue.

(28) Non-compliance with the deployer obligations, including Article 27, is subject to administrative fines of up to EUR 15 000 000 or, for an undertaking, up to 3 % of total worldwide annual turnover, whichever is higher, under Article 99(4) AI Act.

Annex

Template for the Fundamental Rights Impact Assessment required by Article 27 of Regulation (EU) 2024/1689 (AI Act)

Version of the assessment	*Version identifier, with references to previous versions where applicable.*
Last update	*DD/MM/YYYY*
Performed by	*Name, role and function with responsibility for the assessment.*
Contributors	*Each function that contributed and in what capacity — process owner, technical, legal, data protection, representatives of affected persons.*
Approved by	*Name and role of the person approving deployment on the basis of this assessment.*
Trigger determination reference	*Reference and date of the completed ECS-AIA-HR-SCOPE establishing that the Article 27 duty applies.*
Classification	*Internal / disclosable to authorities / publishable in whole or in part.*

1. System and deployment

Deployer	*Legal name and, where relevant, the part of the organisation deploying the system.*
Deployer category	*Body governed by public law / private entity providing public services / deployer of an Annex III point 5(b) or 5(c) system.*
System, provider and version	*Identify the system and its provider.*
Annex III classification	*The point and sub-point under which the system is classified.*
Intended purpose per the provider	*As stated in the instructions for use.*
Date of first intended use	*DD/MM/YYYY. This assessment must be completed before this date.*

2. Element (a) — the deployer's processes

The process, end to end	*What happens before the system is invoked, what it produces, who receives the output, what they do with it, and what the affected person experiences. Describe the process, not the system.*
Where the output bears on a person	*Identify the precise point at which the system's output affects a decision, an allocation, or a person's treatment.*
Nature of the decision	*Fully automated / output followed in substantially all cases / output one input among several / output used for triage or prioritisation only. Where the practical effect differs from the formal position, record the practical effect.*
Consistency with intended purpose	*Confirm the use is in line with the intended purpose stated by the provider, and describe any divergence.*
What happens if the system is unavailable	*The fallback process, and whether it is capable of operating at the volumes concerned.*

3. Element (b) — period and frequency of use

Period of intended use	*From when, and for how long. Identify any review or reauthorisation point.*
Frequency and volume	*How often the system is used and the approximate number of persons affected in a defined period.*
Phasing	*Whether deployment is phased or piloted, the stages, and the criteria for progressing.*

4. Element (c) — affected categories and groups

Identify sub-groups likely to experience different outcomes. Element (d) requires risks to be assessed against the categories identified here, so a population described only in general terms cannot support the analysis.

Category or group	How they are affected	Direct or indirect	Why this group is identified separately

Groups in a position of vulnerability	*Identify groups with limited capacity to detect or contest an adverse outcome, or for whom an adverse outcome has consequences of a different order.*
Under-representation	*Groups under-represented in the data on which the system was developed, or for whom its inputs are less reliable, so far as the provider's information discloses this.*

Indirectly affected persons	*Persons affected without being processed — dependants, household members, others affected by an allocation decision.*

5. Element (d) — specific risks of harm

Express risks as harms to people, not failures of the system. Differentiate across the groups identified in Section 4. Take into account the information given by the provider under Article 13.

Provider information relied upon	*Identify the Article 13 information used — accuracy metrics, known limitations, performance across populations, conditions of intended use. Record where the provider's information is insufficient to assess a risk and what was asked.*

Risk of harm	Group(s) affected	Rights engaged	Likelihood and severity	Source of the risk

Discrimination analysis	*How differential outcomes across protected characteristics were assessed, including consideration of proxies. Where no analysis was possible, state why and what was requested from the provider.*
Risks from correct operation	*Risks arising where the system operates exactly as designed, as distinct from risks of error.*
Effective remedy	*Whether an affected person could realistically understand, detect and contest an adverse outcome, assessed from their position rather than the organisation's.*
Cumulative and systemic effects	*Effects arising from scale, repetition, or interaction with other systems and processes affecting the same population.*

6. Element (e) — human oversight measures

Who exercises oversight	*Role, function and reporting line. Cross-refer to ECS-AIA-HR-OVERSIGHT.*
Competence, training and authority	*What qualifies them, what training they receive, and what authority they hold to override, escalate or suspend.*
What the overseer can see	*The information available at the point of oversight, including whether reasons for the output are available.*
What triggers intervention	*The conditions under which the overseer is expected to depart from or question the output.*
Practical capacity	*Time available per case, caseload, and any operational pressure bearing on the ability to exercise oversight meaningfully.*
Risks oversight mitigates	*Which of the risks in Section 5 oversight is expected to address.*
Risks oversight cannot mitigate	*Which risks oversight cannot address, stated expressly. This is the more important of the two rows.*

7. Element (f) — measures on materialisation, governance and complaints

Measures should be specific, owned and triggered: what happens, who does it, and what sets it in motion.

Risk (from Section 5)	Measure if it materialises	Trigger condition	Owner

Internal governance arrangements	*Who owns the deployment, who reviews its operation, at what frequency, and to whom escalation runs.*
Suspension arrangements	*The conditions under which use would be suspended, who decides, and how quickly it can be effected.*
Complaint mechanism	*How an affected person raises a complaint about a decision involving the system. Describe it from their position.*
How the person learns the system was involved	*The information given to affected persons, and where the Article 26(11) duty is discharged.*
Effect of a complaint on the decision	*What happens to the decision while a complaint is considered.*
Who resolves complaints	*The function responsible, its independence from the process owner, and target timescales.*
Accessibility of the complaint route	*How persons who may struggle with a standard route are accommodated.*

Redress where harm has occurred	*What is done for a person who has been adversely affected, beyond correcting the decision.*

8. Data protection impact assessment interface

DPIA carried out?	*Yes / no, with reference and date.*
Elements met through the DPIA	*Identify which of elements (a) to (f) are addressed in the DPIA and where, so that this assessment complements rather than duplicates it.*
Elements addressed only here	*Identify the elements this assessment addresses that the DPIA does not.*
Data protection officer involvement	*Whether the DPO was consulted and when.*

9. Conclusion and deployment decision

Overall conclusion	*Whether the deployment should proceed, proceed subject to conditions, be modified, or not proceed. State which.*
Conditions of deployment	*Any condition on which the decision to deploy depends, its owner and its deadline.*
Residual risks accepted	*Risks that remain after mitigation and are accepted, identified expressly, together with who accepted them.*
Limitations of this assessment	*What could not be assessed, and why. Include any absence of provider information and any absence of input from affected persons.*
Alternatives considered	*Whether the purpose could be achieved without this system, and why this deployment was preferred.*

10. Notification to the market surveillance authority

Authority notified	*The market surveillance authority notified, and the date.*
Form of notification	*Whether the template or automated tool referred to in Article 27(5) was used, and its version. Where it was not available, describe what was submitted.*
Results notified	*Summarise what was communicated as the results of the assessment.*
Any exemption relied upon	*Where an exemption from the notification obligation is relied upon, identify its basis.*
Correspondence held	*Where the notification and any response are retained.*

11. Reliance on prior assessments

Prior assessment relied upon	*Identify any previously conducted fundamental rights impact assessment or provider assessment relied upon.*
Basis for similarity	*Why the cases are similar — addressing the process, the affected population and the scale.*
What has been assessed afresh	*The elements assessed independently notwithstanding the reliance.*

12. Review and updating

Review cadence	*How often the assessment is reviewed in the absence of a trigger.*
Trigger events registered	*The changes that would render an element no longer up to date — changes to the process, the population, the volume, the system, the oversight arrangements, or the provider's stated performance.*

How triggers are detected	*Who monitors for them, and how.*
Last review	*DD/MM/YYYY and the outcome.*

13. Version history and declaration

Version	Date	Nature of the change or review	Approved by

Declaration

This assessment has been performed prior to the first use of the high-risk AI system to which it relates. The information recorded is provided in good faith and, to the best of the knowledge of the person signing, is accurate and complete as at the date of signature. Risks and limitations have been recorded as found, including those that could not be mitigated.

Name	
Role	
Date	
Signature	