AI GOVERNANCE TEMPLATE

AI Impact Assessment

A defensible, framework-aligned template for documenting the risk and impact of any AI system you build, buy, or deploy.

EU AI Act	NIST AI RMF	ISO/IEC 42001	Colorado AI Act	TRAIGA

VERSION
1.0

PUBLISHED
April 2026

Published by aicompliancevendors.com — an independent buyer's guide for AI governance, risk and compliance tooling. This template is editorial, vendor-neutral, and ungated.

INTRODUCTION

How to use this template

An AI Impact Assessment (AIIA) is the document a regulator, customer, or internal risk committee will ask for first. It records what the system does, who could be harmed, what mitigations are in place, and who is accountable. Done well, it is the central artefact in your AI governance programme — referenced by your risk register, your vendor due diligence file, and your board reporting.

This template is structured into eight sections and one mapping table. Complete one assessment per AI system, where "system" means a discrete AI capability with a single intended use. If the same model is used for two distinct purposes (e.g. a hiring screen and a customer support assistant), produce two assessments.

WORKFLOW

1. Scope it.	Section 1 fixes the boundary of the assessment. Lock the system description and intended use before completing any other section.
2. Identify who is affected.	Section 2 forces explicit identification of stakeholders, including third-parties whose data flows through the system.
3. Surface risk before mitigation.	Sections 3 and 4 should be completed before considering controls. The unmitigated picture is what regulators ask for.
4. Document fairness empirically.	Section 5 requires data, not assertions. Where data does not yet exist, document the gap and the plan to close it.
5. Tie controls to risks.	Every control in Section 6 should reference a specific risk from Section 4. If it does not, it is decoration.
6. Define oversight in operational terms.	Section 7 must describe what a human will actually do, with what evidence, on what cadence.
7. Set update triggers.	Section 8 specifies the events that require this assessment to be revisited. AIIAs are living documents.
8. Map to your obligations.	Section 9 is the regulator-facing crosswalk. Update only the columns that apply to your jurisdiction.

SECTION 1

System description & intended use

Anchor the assessment in a single, unambiguous description of the system and its purpose. This section is the boundary that every later section relies on.

SYSTEM NAME

Internal product or workstream name.

SYSTEM OWNER (ACCOUNTABLE EXECUTIVE)

Named individual, role, and reporting line.

AI PROVIDER / VENDOR (IF EXTERNAL)

Vendor name, model name, version, and contract reference.

INTENDED USE

What is this system used for? What problem does it solve? Be specific. “Generative AI assistant for customers” is not specific. “Drafts initial responses to billing queries for human review by support agents” is.

SECTION 1 (CONTINUED)

System description & intended use

OUT-OF-SCOPE USES

Uses that this assessment does NOT cover. Include uses that are explicitly prohibited and any planned future use that requires a separate AIIA.

MODEL / SYSTEM ARCHITECTURE SUMMARY

Foundation model, fine-tuning, retrieval-augmented generation, agents, deterministic logic — describe at a level a non-engineer can follow.

DATA SOURCES USED AT TRAINING, FINE-TUNING, AND INFERENCE

List datasets, vendors, and customer data flows. Note any sensitive categories (health, biometric, children, protected characteristics).

DECISION AUTHORITY OF THE AI

Does the system recommend, decide, or act? Is a human always in the loop? Is the human reviewing all outputs, a sample, or only escalations?

SECTION 2

Stakeholder identification

Identify everyone whose interests are affected by the system, including people who never interact with it directly. The categories below come from the EU AI Act's "affected persons" concept and the NIST AI RMF Map function.

STAKEHOLDER GROUP	WHO SPECIFICALLY	HOW THEY ARE AFFECTED	INPUT (Y/N)
End users / customers			
Subjects of decisions (not direct users)			
Employees who operate the system			
Third parties whose data is processed			
Regulators with jurisdiction			
Communities / public at large			
Other (specify)			

SECTION 3

Risk identification

Use the four NIST AI RMF risk categories below as a check. For each category, list the specific risks this system poses. A category may have no risks — but if so, document why.

TECHNICAL RISK

Model performance failures, drift, hallucination, security vulnerabilities, brittle behavior on edge cases.

OPERATIONAL RISK

Failures in workflow integration, ambiguous handoffs to humans, dependency on a single vendor, costs that scale unpredictably.

LEGAL & REGULATORY RISK

Non-compliance with EU AI Act, NIST, ISO 42001, Colorado AI Act, TRAIGA, GDPR, sector-specific rules, employment law.

REPUTATIONAL & ETHICAL RISK

Harm to specific groups, undermining of user trust, environmental cost of training, public perception of misuse.

SECTION 4

Risk categorisation & prioritisation

Score each risk identified in Section 3 on severity and likelihood. Severity considers the magnitude of harm to affected stakeholders. Likelihood reflects pre-mitigation probability. The product is the inherent risk score — the number you carry into Section 6.

SCALE	1 — Minimal	2 — Low	3 — Moderate	4 — High	5 — Severe
Severity	Negligible impact, easily reversed	Limited impact on individuals	Material impact, recoverable	Significant or hard to reverse	Severe, irreversible, or systemic
Likelihood	Rare	Unlikely	Possible	Likely	Almost certain

#	RISK	STAKEHOLDER	SEVERITY	LIKELIHOOD	SCORE
1					
2					
3					
4					
5					
6					
7					

SECTION 5

Fairness & bias assessment

Assertions are not assessments. This section requires data on protected characteristics where possible and a documented gap-closure plan where not. The fields below align to the metrics expected by the EU AI Act Article 10 (data and data governance), NIST AI RMF Measure function, and NYC Local Law 144 four-fifths-rule disclosures.

PROTECTED CHARACTERISTICS IN SCOPE

Race, gender, age, disability, religion, sexual orientation, pregnancy, national origin, veteran status — mark each that is in scope and why.

REPRESENTATIVENESS OF TRAINING & EVALUATION DATA

Compare distribution of protected groups in your data against the population the system will serve. Note material gaps.

BIAS METRICS MEASURED

List metrics: e.g. demographic parity, equalised odds, four-fifths impact ratio, calibration by group, false positive rate parity. State current values.

DISPARATE IMPACT ANALYSIS RESULT

If using the four-fifths rule (NYC LL 144 standard): record the impact ratio for each protected group versus the most-selected group.

ACCOMMODATIONS & ALTERNATIVES

Where law requires (e.g. NYC LL 144, Illinois HB 3773), describe the alternative process available to candidates or affected persons.

KNOWN LIMITATIONS & GAPS

Be explicit. “We do not yet have data on disability status” is a defensible position; silence is not.

SECTION 6

Mitigation measures & controls

For each risk identified in Section 4, document the controls in place and the residual score after the control is applied. A control without a linked risk is not a control — it is decoration.

RISK #	CONTROL DESCRIPTION	TYPE	OWNER	RESIDUAL SCORE	EVIDENCE REFERENCE

Control type guide: Preventive stops the risk from occurring (e.g. input validation). Detective finds it when it happens (e.g. drift monitor). Corrective reverses the impact (e.g. rollback procedure). Compensating reduces likelihood or severity but doesn't directly address the risk (e.g. insurance, manual review).

SECTION 7

Human oversight provisions

Required by EU AI Act Article 14, expected by NIST AI RMF, and a central element of state AI laws including Colorado SB 24-205 and the TRAIGA disclosure framework. Describe what humans actually do — with what evidence, on what cadence, with what authority to override.

OVERSIGHT MODEL

Human-in-the-loop (review every output), human-on-the-loop (sample / supervise), human-out-of-the-loop (post-hoc audit only). State which and why.

REVIEWER ROLE & COMPETENCE

Job title and required training of the human(s) performing oversight. Cite the training programme and refresh cadence.

WHAT THE REVIEWER SEES

The exact information presented at the point of decision: input, model output, confidence score, explanation, recommended action.

OVERRIDE AUTHORITY & PROTOCOL

How a human overrides the system, when override is mandatory, how overrides are logged, who reviews override patterns.

ESCALATION PATH

When oversight reveals a serious problem, who is paged, on what timeline, with what authority to suspend the system.

SECTION 8

Monitoring & update triggers

An AIIA is a living document. The events below should trigger a documented re-assessment. Any “Yes” in the right column means this assessment must be updated and re-approved.

The model is retrained, fine-tuned, or replaced.	☐
The intended use changes or expands.	☐
A new stakeholder group becomes affected.	☐
A serious incident is recorded (per EU AI Act Art. 73 / your incident policy).	☐
Bias monitoring detects movement outside accepted tolerance.	☐
Drift, hallucination rate, or refusal rate moves outside accepted tolerance.	☐
A new regulation, guidance, or enforcement action affects the system’s classification.	☐
The vendor or provider materially changes terms, ownership, or model architecture.	☐
Twelve months have elapsed since the last review (mandatory annual refresh).	☐

ASSESSMENT SIGN-OFF

Author	Reviewer	Approver
Name & role	Name & role	Name & role
Date	Date	Date

SECTION 9 · CROSSWALK

Mapping to regulatory frameworks

Each row of this template is referenced by one or more frameworks. Use this mapping to cross-reference your assessment against the obligations that apply to your organisation. Update only the columns for jurisdictions in scope.

SECTION	EU AI ACT	NIST AI RMF	ISO/IEC 42001	COLORADO AI ACT	TRAIGA
1. System description	Art. 11 + Annex IV §§1–3	Map 1–4	Annex A.6.2.2 / 7.5	§ 6-1-1701 (defs.)	§ 552.001 (defs.)
2. Stakeholders	Art. 27 (FRIA)	Map 5	Annex A.7.4 / 8.4	§ 6-1-1703 risk-mgmt	§ 552.052(b)
3. Risk identification	Art. 9 §§1–4	Map 5; Measure 1–2	6.1 / Annex A.5	§ 6-1-1703(2)	§ 552.054(b)
4. Risk categorisation	Art. 9 §5; Annex III	Manage 1–4	6.1.2 / Annex A.5	§ 6-1-1703(2)(a)	§ 552.054(b)(2)
5. Fairness & bias	Art. 10 §§2–5	Measure 2.11; 3	Annex A.6.2.4 / 8.5	§ 6-1-1703(2)(b)	§ 552.052(c)
6. Mitigations & controls	Art. 9 §§3–5; Art. 15	Manage 1.4 / 2.1	Annex A.5–A.10	§ 6-1-1703(3)	§ 552.052(b)(2)
7. Human oversight	Art. 14	Govern 4; Manage 1.3	Annex A.6.2.6 / 8.6	§ 6-1-1703(4)	§ 552.054(c)
8. Monitoring & updates	Art. 17 / 72; Art. 73	Manage 4; Measure 4	9.1 / 10.1–10.2	§ 6-1-1703(5)	§ 552.054(b)(4)

Citations are to publicly available statutory text and authoritative published guidance: EU AI Act Regulation (EU) 2024/1689 (Articles 9–17, 27, 72–73, Annex IV); NIST AI RMF 1.0 (January 2023); ISO/IEC 42001:2023 AI Management Systems — Annex A controls; Colorado AI Act (SB 24-205) Colo. Rev. Stat. § 6-1-1701 et seq.; TRAIGA Texas Responsible AI Governance Act (HB 149) Tex. Bus. & Com. Code Ch. 552.

APPENDIX

Sources & further reading

This template draws exclusively on official statutory text and primary-source standards. Every framework reference in Section 9 is verifiable against the URLs below.

1. EU AI Act — Regulation (EU) 2024/1689
   https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689

2. NIST AI Risk Management Framework 1.0 (January 2023)
   https://www.nist.gov/system/files/documents/2023/01/26/AI%20RMF%201.0.pdf

3. ISO/IEC 42001:2023 — AI Management Systems
   https://www.iso.org/standard/81230.html

4. Colorado AI Act (SB 24-205) — Colo. Rev. Stat. § 6-1-1701 et seq.
   https://leg.colorado.gov/sites/default/files/2024a_205_signed.pdf

5. Texas TRAIGA (HB 149) — Tex. Bus. & Com. Code Ch. 552
   https://capitol.texas.gov/tlodocs/89R/billtext/pdf/HB00149F.pdf

6. NYC Local Law 144 — Automated Employment Decision Tools
   https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page

7. Illinois HB 3773 — Human Rights Act AI Amendment
   https://www.ilga.gov/legislation/publicacts/103/103-0804.htm

8. Companion guide on aicompliancevendors.com
   https://aicompliancevendors.com/blog/ai-impact-assessment-template

About aicompliancevendors.com

An independent buyer's guide for AI governance, risk and compliance tooling. We publish vendor-neutral frameworks, comparison guides, and downloadable templates under an open-access editorial policy. Every claim in our content is traceable to a public source. We do not accept paid vendor placements for editorial coverage.

aicompliancevendors.com