

AI Governance Tool Selection Checklist

Used internally by governance, compliance, and AI risk teams

Selecting AI governance software
aligned with
EU AI Act compliance requirements



AI Governance Desk

Internal & Advisory Use

Version 1.1

How to Use This Checklist



This checklist is designed to help organizations evaluate and select AI governance software aligned with the European Union AI Act and emerging global AI regulatory expectations.

It provides a structured, technology-agnostic framework for assessing whether governance platforms support key compliance pillars such as data traceability, model monitoring, risk management, and audit readiness.

This checklist does not recommend specific vendors. Instead, it enables informed, defensible decision-making across governance, compliance, and technical teams.

Intended Audience

This checklist is intended for:

- AI governance and oversight teams
- Compliance and legal professionals
- Chief Risk Officers (CROs)
- Chief Data Officers (CDOs)
- Machine learning and AI engineering leaders
- Internal audit and assurance teams

Scoring Method

Each item should be scored using the following scale:

- Yes — The capability is fully supported and operational
- Partial — The capability exists but is limited, manual, or incomplete
- No — The capability is not supported

Scoring should reflect current functionality, not roadmap promises.

How to Apply the Checklist

1. Review each section carefully before scoring.
2. Each section corresponds to a critical AI governance capability required for compliant AI deployment.
3. Score each criterion honestly based on evidence.
4. Use documentation, platform demonstrations, and technical reviews—avoid assumptions.
5. Use the scoring guidance consistently.
6. Apply the same standards across all tools being evaluated to ensure fairness and comparability.
7. Identify gaps early.
8. Low or partial scores highlight governance risks that may require remediation, additional controls, or alternative tooling.
9. Retain completed checklists for audit purposes.
10. Completed versions can serve as internal decision records and support future regulatory assessments.

This checklist is provided for informational and internal evaluation purposes only. It does not constitute legal advice, regulatory approval, or certification of compliance under the EU AI Act or any other law.

Organizations remain responsible for ensuring compliance with applicable regulations and should consult qualified legal or regulatory experts where appropriate.

Governance Capability Areas



This checklist evaluates AI governance software across three foundational capability areas that collectively support compliant, scalable, and defensible AI deployment.

These capability areas reflect the practical controls regulators expect organizations to demonstrate under the EU AI Act and emerging global AI governance frameworks. While tools may emphasize different features, effective governance requires coverage across all three areas.

four foundational capability areas

Capability 1 – Data Lineage & Observability

This capability focuses on end-to-end visibility into how data is sourced, transformed, and used throughout the AI lifecycle.

Strong data lineage and observability enable organizations to demonstrate traceability, data quality controls, bias management, and version awareness – all of which are central to Article 10 data governance obligations for high-risk AI systems.

Key outcomes supported:

- Traceable data provenance
- Bias and data quality visibility
- Versioned datasets and features
- Audit-ready data flows

Capability 2 – Model Monitoring & Drift Detection

This capability addresses post-deployment oversight of AI systems, ensuring that models remain reliable, fair, and robust over time.

Continuous monitoring supports detection of performance degradation, data drift, concept drift, and emerging bias risks. These controls are essential for maintaining compliance with robustness and post-market monitoring expectations under the EU AI Act.

Key outcomes supported:

- Performance tracking over time
- Drift and anomaly detection
- Fairness and stability monitoring
- Early identification of emerging risks

Capability 3 – Documentation & Audit Management

This capability ensures that governance evidence is systematically captured, maintained, and exportable for internal review, audits, or regulatory inquiries.

Documentation and audit management enable organizations to assemble structured technical files, maintain decision records, and demonstrate conformity without relying on ad-hoc or manual processes.

Key outcomes supported:

- Structured governance documentation
- Version-controlled evidence
- Annex IV-aligned technical files
- Exportable audit packages

Capability 4 – AI Decision Traceability & Audit Trails

This capability focuses on reconstructing and explaining individual AI system decisions in context, including inputs, system behavior, outputs, and oversight actions.

It represents an emerging but critical layer of AI governance, supporting transparency, accountability, and auditability under evolving regulatory expectations.

Transition to Scoring Sections

While some platforms may offer partial coverage across these areas, effective AI governance requires deliberate evaluation of each capability independently.

Scoring Framework & Rating Scale



This checklist uses a standardized scoring framework to support objective, evidence-based evaluation of AI governance platforms.

Each capability within the checklist should be assessed based on current, demonstrable functionality, not future roadmap commitments or marketing claims.

Scoring should be supported by documentation, live demonstrations, technical validation, or operational evidence wherever possible.

Scoring Scale

Each item should be scored using the following scale:

Yes

The capability is fully supported, operational, and available without significant manual workarounds.

Partial

The capability exists but is limited, requires manual processes, custom development, or does not fully meet governance requirements.

No

The capability is not supported or is materially insufficient for compliant AI governance use.

Scoring Guidance

- Score conservatively and consistently across all tools evaluated
- Apply the same standards to every vendor or platform
- Avoid assumptions or inferred capabilities
- Roadmap features should not influence scoring
- Partial scores should be documented with rationale

Low or partial scores may indicate governance gaps that require remediation, compensating controls, or alternative tooling.

Use of Results

Completed checklists may be used to:

- Support internal procurement decisions
- Identify governance risk gaps
- Inform remediation planning
- Provide evidence during audits or regulatory inquiries
- Maintain defensible records of AI governance decisions

Important Note

This checklist supports internal evaluation and decision-making. It does not constitute legal advice, regulatory approval, or certification of compliance under the EU AI Act or any other law.

Organizations remain responsible for ensuring compliance with applicable regulations and should consult qualified legal or regulatory experts where appropriate.

Capability 1 – Data Lineage & Observability

This section evaluates whether governance platforms provide end-to-end visibility into how data is sourced, transformed, versioned, and used throughout the AI lifecycle.

Strong data lineage and observability are foundational requirements under Article 10 of the EU AI Act.

Organizations must be able to demonstrate where training and operational data originated, how it was processed, which versions were used, and how risks such as bias and data quality issues are managed over time.

Capability Area	Description	Score (Yes / Partial / No)	Notes / Evidence
Data Source Traceability	Ability to identify, document, and trace all data sources used for training, validation, and inference, including origin, ownership, and usage rights.		
Feature Engineering Lineage	Visibility into feature creation, transformations, and dependencies, enabling reconstruction of how model inputs were derived.		
Dataset Versioning	Support for version-controlled datasets and features, including historical snapshots linked to specific model versions.		
Bias & Data Quality Monitoring	Mechanisms to assess and document bias, data quality issues, and distributional changes across datasets over time.		
Audit-Ready Lineage Records	Ability to export lineage documentation and evidence suitable for internal audits or regulatory review.		

Scoring Guidance

Score each item based on current, operational capabilities.
 Partial scores should be supported with documented limitations or manual dependencies.

Capability 2 – Model Monitoring & Drift Detection



This section evaluates whether governance platforms provide continuous post-deployment oversight of AI systems to ensure models remain reliable, fair, and robust over time.

Under the EU AI Act, organizations are required to monitor high-risk AI systems throughout their operational lifecycle. This includes detecting performance degradation, data drift, concept drift, emerging bias, and other systemic risks that may arise after deployment.

Effective model monitoring enables early intervention, supports post-market monitoring obligations, and provides evidence that AI systems continue to operate within acceptable risk thresholds.

Capability Area	Description	Score (Yes / Partial / No)	Notes / Evidence
Performance Monitoring Over Time	Ability to track model accuracy, error rates, confidence scores, and key performance indicators continuously after deployment.		
Data Drift Detection	Mechanisms to detect shifts in input data distributions that may affect model reliability or validity.		
Concept Drift Detection	Ability to identify changes in the underlying relationship between inputs and outputs that degrade model performance over time.		
Bias & Fairness Monitoring	Ongoing monitoring of fairness metrics across protected or relevant population groups to identify emerging bias.		
Alerting & Threshold Management	Configurable alerts and thresholds that notify stakeholders when performance, drift, or fairness metrics exceed acceptable limits.		
Audit-Ready Monitoring Records	Ability to retain logs, monitoring outputs, alerts, and remediation actions as evidence for internal audits or regulatory review.		

Scoring Guidance

Score each item based on current, operational monitoring capabilities.

- Is manual
- Is periodic rather than continuous
- Lacks alerting or documentation
- Does not cover fairness or drift comprehensively

Partial scores should be assigned where monitoring exists but:

Document limitations, assumptions, and remediation plans in the Notes / Evidence column.

Capability 3 – Documentation & Audit Management



This section evaluates whether governance platforms systematically capture, manage, and export documentation required to demonstrate AI governance compliance during internal reviews, audits, or regulatory inquiries.

Under the EU AI Act, organizations deploying high-risk AI systems must maintain comprehensive technical documentation, decision records, and conformity evidence throughout the AI lifecycle.

Effective documentation and audit management reduce reliance on ad-hoc processes and ensure that compliance evidence is complete, traceable, and audit-ready at any point in time.

Capability Area	Description	Score (Yes / Partial / No)	Notes / Evidence
Technical Documentation Management	Ability to create, maintain, and update structured technical documentation aligned with EU AI Act Annex IV requirements.		
Model Cards & System Documentation	Support for standardized model cards, system descriptions, intended use statements, and known limitations.		
Risk Assessment Records	Centralized storage and version control of AI risk assessments, mitigation actions, and governance approvals.		
Decision & Change Logs	Ability to record governance decisions, approvals, overrides, and changes across the AI lifecycle.		
Audit Trail & Evidence Retention	Retention of audit-ready logs and evidence supporting conformity assessments, internal audits, or regulator requests.		
Exportable Audit Packages	Ability to export complete documentation bundles suitable for notified bodies, regulators, or internal assurance teams.		

Scoring Guidance

Score each item based on current documentation and audit capabilities.

Partial scores should be assigned where documentation exists but is fragmented, manually maintained, or not consistently version-controlled.

Notes should clearly indicate documentation gaps, manual dependencies, or remediation actions.

Capability 4 – AI Decision Traceability & Audit Trails



This capability evaluates whether governance platforms can capture, reconstruct, and explain individual AI system decisions in a structured, audit-ready format.

As AI systems become more autonomous and enforcement under the EU AI Act accelerates, regulators increasingly expect organizations to demonstrate not only how systems are built and monitored, but how specific decisions are generated in context.

Decision traceability extends beyond traditional logging by enabling organizations to reconstruct inputs, reasoning pathways, outputs, and human oversight actions for any given AI decision.

This capability supports transparency obligations under Article 13, strengthens post-market monitoring under Article 72, and enables defensible responses during audits, investigations, or incident reviews.

Capability Area	Description
Decision Input Context	Ability to capture and retain inputs, features, and contextual data used at the time of decision
Model & System Attribution	Identification of model version, pipeline, or system responsible for the output
Decision Reasoning Trace	Visibility into intermediate steps, scoring logic, or reasoning pathways where applicable
Confidence & Outcome Metadata	Recording confidence scores, alternative outputs, or evaluation metrics
Human Oversight Records	Logging approvals, overrides, escalations, and human intervention
Audit-Ready Decision Logs	Ability to export structured decision records for audits and regulatory review

Score each item based on the platform’s ability to provide structured, verifiable decision-level traceability.

Partial scores should be assigned where:

- Decision data is captured but not structured
- Logs exist but are not easily reconstructable
- Reasoning pathways are not accessible
- Human oversight is not consistently recorded

Effective decision traceability should enable reconstruction of any AI output in context, including inputs, system behavior, and oversight actions.

Overall Scoring Summary



This section consolidates scoring results across all three AI governance capability areas to support an overall assessment of governance readiness.

The summary is intended to help decision-makers identify strengths, gaps, and priority areas for remediation when selecting or evaluating AI governance platforms under the EU AI Act and related regulatory frameworks.

Capability Area	Score Summary	Risk Level	Notes
Data Lineage & Observability	Yes	Low	
	Partial	Medium	
	No	High	
Model Monitoring & Drift Detection	Yes	Low	
	Partial	Medium	
	No	High	
Documentation & Audit Management	Yes	Low	
	Partial	Medium	
	No	High	
AI Decision Traceability & Audit Trails	Yes	Low	
	Partial	Medium	
	No	High	

Overall Governance Readiness Assessment

- Strong governance readiness
- Moderate governance readiness (remediation required)
- High governance risk (not suitable for regulated AI use)

This assessment reflects the lowest-scoring capability area, as gaps in any core governance function may introduce compliance, operational, or regulatory risk.

How to Interpret Results

- Strong readiness indicates broad, operational support across all governance areas.
- Moderate readiness suggests partial coverage with identifiable remediation needs.
- High governance risk indicates material gaps that may prevent compliant AI deployment under the EU AI Act.

Organizations should avoid selecting platforms with critical gaps in any capability area supporting high-risk AI systems.

Interpretation & Decision Guidance



How to translate checklist results into defensible AI governance decisions

This page provides guidance on how to interpret checklist results and translate scores into informed, defensible decisions when selecting or evaluating AI governance platforms.

The guidance is designed to support governance, compliance, risk, and technical stakeholders in determining whether a platform is suitable for regulated or high-risk AI use under the EU AI Act and related frameworks.

Decision Guidance by Capability Outcome

If All Capability Areas Score “Yes”

- Platform demonstrates strong AI governance readiness
- Suitable for regulated and high-risk AI deployments
- Proceed with procurement or onboarding, subject to standard security and contractual reviews

If One or More Areas Score “Partial”

- Platform shows partial governance readiness
- Gaps must be clearly documented and remediated
- Use only where compensating controls or remediation plans are feasible
- Not recommended for high-risk AI use without mitigation

If Any Capability Area Scores “No”

- Platform presents material AI governance risk
- Not suitable for regulated or high-risk AI systems
- Selection should be avoided unless governance gaps are resolved prior to deployment

Critical Rule for Decision-Making

Governance readiness is determined by the lowest-scoring capability area.

A single critical gap in data lineage, model monitoring, or documentation may introduce regulatory, operational, or audit risk, regardless of strengths in other areas.

Important Note

This checklist supports internal evaluation and governance decision-making.

It does not certify compliance, replace regulatory interpretation, or constitute legal advice. Organizations remain responsible for ensuring compliance with applicable laws and regulatory obligations.

Recommended Actions Based on Results

- Document all Partial and No scores with supporting evidence
- Identify whether gaps can be addressed through tooling, process, or controls
- Escalate high-risk gaps to governance or risk committees
- Retain completed checklists as part of procurement and audit records



AI Governance Desk

AI Governance Desk provides practical governance tools, compliance frameworks, and readiness assessments to support organizations deploying or managing AI systems under emerging global AI regulations, including the EU AI Act.

Our resources are designed to support internal governance, risk management, compliance planning, and executive decision-making across high-impact and high-risk AI use cases.

Why This Checklist Matters

- Support objective evaluation of AI governance platforms
- Identify governance gaps before deployment or procurement
- Enable defensible, evidence-based decision-making
- Support audit readiness and regulatory inquiries
- Align internal governance with EU AI Act expectations

Intended Use

This checklist is intended for internal governance, compliance, risk assessment, and advisory use. It is designed to support structured evaluation and internal decision-making, not to certify compliance or replace regulatory or legal review.

DISCLAIMER

This document does not constitute legal advice, regulatory approval, or certification of compliance under the EU AI Act or any other law.

Organizations remain responsible for ensuring compliance with applicable regulations and should consult qualified legal, regulatory, or professional advisors where appropriate.