

EU AI Act Compliance Report

eu_ai_act_compliance

Generated 14 July 2026 · report id sr-5470619c8a4b · \$1.50 USDC on Base via x402

SAMPLE REPORT · REAL TOOL EXECUTION AGAINST A FICTIONAL INPUT · A PAID CALL RETURNS THE FULL REPORT FOR YOUR OWN INPUT

high_risk

RISK TIER

5

OBLIGATION GAPS

3

CRITICAL GAPS

2(a)

recruitment

ANNEX III BASIS

PROVENANCE & PAYMENT			
PAYMENT		EXECUTION & INTEGRITY	
Price	\$1.50 USDC	Report id	sr-5470619c8a4b
Network	Base (chain id 8453)	Artifact id	5470619c-8a4b-467b-8dbb-e87d382053f3
Protocol	x402 v2, scheme exact	Execution id	48d7aa0b-80bf-4f34-b2da-a9db726c68fb
Pay to	0x97a1...D9a4	Tool	eu_ai_act_compliance v1.0.0
Payer	0x72f41a8f59e92e553ed6... (fictional)	Schema	Output conforms to the published JSON schema for eu_ai_act_compliance.
Auth nonce	0x2ce877204476d40d10794a74... (fictional)	Input sha-256	38a70275a882473d279412a8a78f3a40... (sha-256 of the exact tool input)
Facilitator	✓ verified ✓ settled	Generated	14 July 2026
CHAIN		Endpoint	https://systemprompt.io/api/v1/mcp/systemprompt-reports/mcp
Settlement tx	0xc3013aa1a622b0120ec9c53fc32209fe... (illustrative hash for this sample)	This sample is the output of a genuine tool execution against a fictional input; the artifact and execution ids above identify the real run. Discovery and validation are free. Payment is only requested after input passes the free validation gate, and invalid input is never charged.	
Explorer	https://basescan.org/tx/<hash>		
USDC contract	0x833589fCD6eDb6E08f4c7C32D4f71b54bdA02913		
Settled at	14 July 2026 14:07 UTC		
A real report's X-PAYMENT-RESPONSE header carries the genuine on-chain settlement hash, verifiable on BaseScan.			

■ VALIDATION CHECKLIST

5 of 5 stages executed on this run.

- ✓

Validate

DETERMINISTIC

Passed

The system description and data-flows note cleared validation: both were above the 50-character minimum and the combined input was well under the 60,000-character cap, so processing continued.

Required fields and minimum sizes are checked before any payment. system_description must contain at least 50 characters of substantive content and the combined input is capped at 60,000 characters. Invalid input is rejected free of charge with a structured...
- ✓

Parse

DETERMINISTIC

Passed

No decision_tree_answers were supplied, so the codified Annex III decision tree was not invoked; the description passed through to the analysis stage without a deterministic classification.

When decision_tree_answers is supplied, the codified Annex III decision tree runs before any charge. It walks Article 3(1) definition screening, Article 5 prohibitions, the GPAI track, the Annex I product route, and the Annex III areas. If the answers are i...
- ✓

Rule engine

DETERMINISTIC

Ran

The high_risk classification was resolved to its triggered obligations, producing gaps against five articles: 9, 10, 13, 14, and 52.
- ✓

Analysis

AI-CONSTRAINED

Ran

The classification pass returned high_risk under Annex III 2(a), and the report pass produced 5 obligation gaps (3 critical, 2 high), a 4-item Annex IV documentation checklist, and a 3-step remediation plan.

Two structured-output passes run today. First, a classification pass by a language model constrained by a fixed JSON schema determines risk tier (prohibited, high_risk, limited_risk, minimal_risk), Annex III categories, prohibited-practice flags, GPAI appli...
- ✓

Render

AI-CONSTRAINED

Ran

The structured JSON report was returned with a rendered markdown artifact.

The structured JSON report is returned together with a rendered markdown artifact in the MCP response, with the deterministic decision-tree analysis attached when it ran.

FINDINGS BY SEVERITY



Input: A fictional description of Northwind Talent Cloud, a SaaS tool that screens and ranks inbound job applications for hiring managers using a model trained on historical hiring decisions, with a data-flows note. No decision-tree answers were supplied, so the classification was reached by the analysis pass rather than the codified tree.

EU AI Act Compliance Report: Northwind Talent Cloud

■ Overview

This report analyzes the Northwind Talent Cloud system against the requirements of the EU AI Act. The system is classified as **HIGH**-risk under **Annex III, 2(a)** (recruitment and selection).

Compliance Gap Analysis

Article	Obligation	Severity
9	Risk Management	CRITICAL
10	Data Governance	CRITICAL
14	Human Oversight	CRITICAL
52	Transparency	HIGH

Limitations & Coverage Notes

- No specific limitations surfaced regarding the system's architecture; however, the lack of explainability (xAI) in the current 0-100 scoring model represents a potential violation of the duty to ensure human control.

Remediation Strategy

Prioritization is set against the August 2026 deadline for **HIGH**-risk systems. Failure to address the lack of human override capabilities (Art 14) and applicant transparency (Art 52) poses the highest risk of enforcement action.

Structured JSON response (excerpt)

The same report is returned as structured JSON conforming to a published schema. Excerpt:

```
{
  "summary": "The Northwind Talent Cloud is a high-risk AI system under Annex III, 2(a) as it facilitates recruitment. Current compliance gaps are significant, specifically regarding transparency, interpretability, and robust human oversight mechanisms.",
  "risk_tier": "high_risk",
  "obligation_gaps": [
    {
      "article": "Article 9",
      "obligation": "Risk Management System",
      "current_state": "No formal risk assessment provided for the lifecycle of the AI system.",
      "gap": "Missing continuous iterative risk management process to identify hazards for candidates.",
      "severity": "critical"
    },
    {
      "article": "Article 10",
      "obligation": "Data Governance",
      "current_state": "Model is trained on historical hiring data.",
      "gap": "Lack of bias detection, mitigation, and dataset representative analysis to ensure fairness in recruitment.",
      "severity": "critical"
    },
    {
      "article": "Article 14",
      "obligation": "Human Oversight",
      "current_state": "No built-in workflow to override or contest rankings.",
      "gap": "Lack of 'human-in-the-loop' controls for overriding AI decisions and preventing automation bias.",
      "severity": "critical"
    },
    {
      "article": "Article 52",
      "obligation": "Transparency for Natural Persons",
      "current_state": "Applicants are not informed of AI usage.",
      "gap": "Non-compliance with the duty to inform candidates that they are subject to an AI-driven recruitment process.",
      "severity": "high"
    }
  ],
  "documentation_checklist": [
    {
      "document": "Risk Management System Documentation",
      "annex_iv_reference": "Section 2",
      "status": "missing"
    },
    {
      "document": "Technical Specifications (Data, Algorithms)",
      "annex_iv_reference": "Section 3",
      "status": "partial"
    },
    {
      "document": "Transparency and Instructions",
      "annex_iv_reference": "Section 5",
      "status": "missing"
    }
  ],
}
```

```
"remediation_plan": [
  {
    "priority": 1,
    "action": "Implement mandatory applicant transparency disclosures (Art 52) and establish a human oversight override mechanism (Art 14).",
    "articles": [
      "Article 14",
      "Article 52"
    ],
    "effort": "large",
    "deadline_context": "Immediate-term to meet transparency and fundamental rights requirements."
  }
],
"disclaimer": "This report is an automated compliance analysis generated for informational purposes only. It does not constitute professional legal advice. Compliance with the EU AI Act requires a holistic review by legal counsel."
}
```

Get the full report

Connect any MCP-capable agent and call `eu_ai_act_compliance` with your own input. Discovery is free; invalid input is rejected before payment.

```
claude mcp add --transport http systemprompt-reports https://systemprompt.io/api/v1/mcp/systemprompt-reports/mcp
```

1. Call the tool; the server replies `402 Payment Required` with x402 terms.
2. Your wallet signs a USDC authorization on Base for \$1.50. No account, no API key.
3. The facilitator verifies and settles the payment on chain.
4. The full report returns in the same call, with the settlement hash in `X-PAYMENT-RESPONSE`.

<https://systemprompt.io/tools/eu-ai-act-compliance-report/>