

EU AI Act

Compliance assessment — 2024/1689

Server: AWS Sage

Slug: aws-sage

Scan id: ddbbe89e0-5350-4e67-8565-5922e50251e2

Assessed at: 2026-05-10 07:05:11 UTC

Sentinel version: 0.4.0

Rules version: 2026-04-23

Non-compliant

DRAFT for review — not legal advice. See attestation block for verification instructions.

Table of contents

1. Executive summary
2. Coverage & transparency
3. Controls summary
4. Control details
5. Multi-step attack chains
6. Cryptographic attestation

1. Executive summary

Assessment of AWS Sage against EU AI Act: overall status non compliant. Of 5 controls, 4 met, 1 unmet, 0 partial, 0 not applicable. 5 control(s) fell within MCP Sentinel's current assessor coverage; remaining control(s) are documented as not_applicable until Phase 6 expands coverage. Unmet controls have findings at or above the framework's mandatory severity threshold and should be remediated before relying on this server in a regulated deployment. All claims are traceable to individual finding rows via finding_id and to the governing rule via rule_id; the enclosing signed envelope commits MCP Sentinel to the exact bytes of this report.

2. Coverage & transparency

Coverage band: low

Coverage ratio: 40%

Rules version: 2026-04-23

Analysis techniques applied:

- ast-taint
- capability-graph
- entropy
- linguistic-scoring
- schema-inference

3. Controls summary

ControlName

Status

Evidence

Art.9Risk

Management

System' Me@

Art.12Record-

Keeping' Me@

Art.13

Transparency

& Provision of Information to Deployers' Unmet1
Art.14Human Oversight' Met 0
Art.15 Accuracy, Robustness, and Cybersecurity' Met0

4. Control details

Art.9 — Risk Management System

22 assessor rule(s) evaluated this control; no findings observed.

' Met

Art.12 — Record-Keeping

5 assessor rule(s) evaluated this control; no findings observed.

' Met

Art.13 — Transparency & Provision of Information to Deployers

14 assessor rule(s) evaluated this control; 1 finding(s) observed (1 critical); at least one finding is at or above the high threshold (status: unmet).

' Unmet

Evidence:

[Critical] F5
(finding 5d391e3f-70d9-43e1-a1d3-866aaa210f4e, confidence 90%)

SOURCE: external-content at initialize.server_name — The MCP client surfaces the server name verbatim in its approval dialog, and the LLM ingests the server name alongside the tool descriptions. A nam

Required mitigations:

- If you own the server and are NOT affiliated with the vendor whose namespace it contains, rename the server to remove the vendor token. Choose a name that makes your actual publisher identity clear. If you ARE a vendor-approved partner and intentionally use the vendor's namespace, request inclusion in the rule's OFFICIAL_NAMESPACES.verified_github_orgs list by publishing the server under a vendor-sanctioned GitHub organisation. Users deciding whether to approve the server should check the repository owner against the vendor's published list of approved partners before granting trust.

Art.14 — Human Oversight

13 assessor rule(s) evaluated this control; no findings observed.

' Met

Art.15 — Accuracy, Robustness, and Cybersecurity

111 assessor rule(s) evaluated this control; no findings observed.

' Met

5. Multi-step attack chains

No multi-step attack chains were synthesized for this server.

6. Cryptographic attestation

Algorithm: HMAC-
SHA256

Key ID: mcp-sentinel-
dev

Signer: mcp-sentinel/
v1

Signed at: 2026-05-14
T08:09:47.063Z

Canonicalization:
RFC8785

HMAC-SHA256 signature (base64, wrapped at 64 chars):

MSGsBE1j5FRATmLiAfbMzilQNhmoTmN4ZWYgKUZVR/c=

Verification instructions:

To verify this report:

1. Extract the report body (everything except the .attestation field).
2. Canonicalize the body via RFC 8785 (JCS).
3. Compute HMAC-SHA256 with the signing key for key_id "mcp-sentinel-dev".
4. Base64-encode the result and compare with the signature above.

