EU AI Act

Compliance assessment — 2024/1689

Server: mcp.deepwiki.com

Slug: mcp-deepwiki-com-20260516082508-a4cfb2
Scan id: ea225ec1-d8ce-44d7-b6db-6ea34d638eb0

Assessed at: 2026-05-16 08:25:08 UTC
Sentinel version: 0.4.0
Rules version: 2026-04-23

! Partially compliant

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

Table of contents

1. Executive summary

Assessment of mcp.deepwiki.com against EU AI Act: overall status partially compliant. Of 5 controls, 4 met, 0 unmet, 1 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. No control is unmet, but partial findings indicate residual risk below the mandatory threshold. 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: medium

Coverage ratio: 70%

Rules version: 2026-04-23

Analysis techniques applied:

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

3. Controls summary

Control	Name	Status	Evidence
Art.9	Risk Management System'	Met	0
Art.12	Record-Keeping'	Met	0
Art.13	Transparency		

& Provision of Information to Deployers' Met 0
Art.14Human Oversight' Met 0
Art.15 Accuracy, Robustness, and Cybersecurity! Partial7

4. Control details

Art.9 — Risk Management System

' Met

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

Art.12 — Record-Keeping

' Met

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

Art.13 — Transparency & Provision of Information to Deployers

' Met

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

Art.14 — Human Oversight

' Met

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

Art.15 — Accuracy, Robustness, and Cybersecurity

! Partial

111 assessor rule(s) evaluated this control; 7 finding(s) observed (7 medium); all findings are below the high threshold (status: partial).

Evidence:

[Medium] B1 (finding 324ac4a6-8623-4c40-8622-fe6a823ae2ba, confidence 77%)
SOURCE: user-parameter at tool read_wiki_structure — Tool "read_wiki_structure" accepts parameters without structural validation. The AI fills each parameter from user input; nothing in the schema rej

[Medium] B1 (finding 48a91f4f-2ee7-4b41-b65e-c0071f3573b6, confidence 77%)
SOURCE: user-parameter at tool read_wiki_contents — Tool "read_wiki_contents" accepts parameters without structural validation. The AI fills each parameter from user input; nothing in the schema rejec

[Medium] B1 (finding 6aec3e33-4bc2-465a-b05e-a7c419b81ef1,

confidence 77%)
SOURCE: user-parameter at tool ask_question — Tool "ask_question" accepts parameters without structural validation. The AI fills each parameter from user input; nothing in the schema rejects injection
[Medium] B6 (finding ee16e871-7791-49c4-a1fa-19b7eeddc96b, confidence 75%)
SOURCE: user-parameter at tool read_wiki_structure — Tool "read_wiki_structure" input_schema accepts arbitrary extra keys. The declared properties are validated, but the handler may read undeclared ke
[Medium] B6 (finding be99f936-40d2-418d-924d-f01b8faaf50b, confidence 75%)
SOURCE: user-parameter at tool read_wiki_contents — Tool "read_wiki_contents" input_schema accepts arbitrary extra keys. The declared properties are validated, but the handler may read undeclared keys
[Medium] B6 (finding f3411b16-ffd9-4940-b8c1-aaad2e03ba37, confidence 75%)
SOURCE: user-parameter at tool ask_question — Tool "ask_question" input_schema accepts arbitrary extra keys. The declared properties are validated, but the handler may read undeclared keys that bypass
[Medium] E1 (finding 5bc7066c-4a00-411f-b6a9-b7325672fd62, confidence 75%)
SOURCE: environment at capability:tools — An MCP server that answers tool enumeration without authentication trusts the network. Under modern threat models (CCS 2007 DNS rebinding, open cloud networki

Required mitigations:

- Add at least one validation keyword to every string and number parameter. For strings: maxLength, pattern, format, or enum. For numbers: minimum, maximum, or multipleOf. JSON Schema validation runs before the tool handler and is the cheapest first-line defence against injection and DoS.
- Set additionalProperties: false on every object schema. This rejects any key outside the declared properties, closing the side-channel smuggling path and enforcing the schema's stated contract.
- Require authentication for all MCP server connections. For remote MCP servers adopt OAuth 2.0 per RFC 9700 / the MCP Authorization specification. For stdio-launched servers rely on the parent process's security boundary and DO NOT expose the same server over network transports. Even localhost-bound servers should require auth: DNS rebinding (CCS 2007) makes localhost reachable from any browser tab.

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: 2026-05-16
T11:10:24.221Z
Canonicalization:
RFC8785

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

```
CpkFWFhBz+U6KxMUxRz0skeeftZpGwZhM9U0uCMtisU=
```

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.
```