

EU AI Act

Compliance assessment — 2024/1689

Server: VERITAS

Slug: veritas

Scan id: 3e6ae1f3-af55-4ed4-a4cf-7eefa94ea699

Assessed at: 2026-05-11 09:30:48 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
2. Coverage & transparency
3. Controls summary
4. Control details
5. Multi-step attack chains
6. Cryptographic attestation

1. Executive summary

Assessment of VERITAS 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: 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' Met
0

Art.14 Human
Oversight' Met
0

Art.15
Accuracy,
Robustness,
and
Cybersecurity!
Partial1

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; no findings observed.

' Met

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; 1 finding(s) observed (1 medium); all findings are below the high threshold (status: partial).

! Partial

Evidence:

[Medium] E1
(finding
95962898-
c74b-4629-ab32-
ffcf7f2c9f18,
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:

- 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: 2026-05-14
T09:28:42.715Z

Canonicalization:
RFC8785

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

MebrmVSjuukZAb6lrftj69FCFmRyRgr4xIoLHDh83a8=

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.

