Article 12 logging support for autonomous AI agents

What MolTrust provides MolTrust is a deterministic cryptographic protocol layer for agent authorization, comparable in role to transport-security and identity layers such as TLS or PKI. Each authorization decision is recorded as a signed entry: the acting agent, the mandate it operated under, and the interaction that produced the result.

- **Agent Authorization Envelope (AAE).** A machine-readable permission contract — explicit scope, rate, and value bounds per agent — enforced before the tool call, not after. Built on W3C DID and Verifiable Credentials.
- **On-chain anchoring on Base L2.** Zero-knowledge proofs and cryptographic hashes of non-personal operational data are anchored on the public ledger. No personal data and no agent payload are written on-chain. A third party can recompute the outcome without access to the recording system.
- **Article 12 logging support.** The evidence produced maps to the EU AI Act Article 12 logging obligations. The transparency obligations under Article 50 apply from 2 August 2026.

What MolTrust does not claim The design supports logging obligations; it does not itself constitute compliance.

- MolTrust does **not** issue compliance certificates. Certification under the EU AI Act, GDPR, or other frameworks is performed by accredited conformity assessment bodies and auditors.
- MolTrust does **not** replace conformity assessment under Article 43. That, and overall compliance responsibility under Article 16, remain with the AI-system provider placing the system on the market.
- MolTrust does **not** provide legal advice. Legal interpretation belongs to qualified counsel.

Why deployment-level evidence matters (LARA) The Aithos LARA evaluation deployed twelve frontier models as workplace agents across more than 3,000 scenarios covering ten EU legal provisions. The best-performing model, Claude Opus 4.7, was lawful in 54% of runs; the worst, Google Gemini 3.1 Pro, in about 10%. Legal compliance in the study ranged from 10% to 54%.

The reading is narrow and empirical: capable models still fail legal compliance under realistic deployment conditions, so compliance has to be enforced as a property of the deployment rather than assumed as a property of the model. The AAE provides that structural property.

Source: Aithos LARA, 12 models, published May 2026 — lara.aithos.org · aithos.org/article/Aithos-LARA.

Standards & references W3C DID Core v1.0 · Verifiable Credentials Data Model v2.0 · IETF Internet-Draft `draft-kroehl-agentic-trust-aae-00` · technical paper (arXiv preprint 2605.06738).

Contact: hello@moltrust.ch · moltrust.ch/compliance.html