

Evidence for what AI agents do.

Why provable records of agent actions become unavoidable between now and 2028 — and what to do about it early.

THE PROBLEM

AI agents have started doing consequential work: paying invoices, approving refunds, moving data, transacting with other agents. After every incident, audit or dispute comes the same question — **what exactly did the agent do, and can you prove it?** Ordinary logs cannot. They are editable, so they are testimony, not evidence.

WHAT SWARRM DOES

Swarm turns every agent action into evidence: a signed receipt in an append-only Merkle log, anchored hourly to a public chain and timestamped by an independent RFC 3161 authority. Any record can be verified by anyone, offline, free — years later, in front of an auditor, a regulator or a court. Sensitive content never leaves your company: the log carries operational metadata and cryptographic fingerprints, never payloads. Adoption is a URL swap or three lines of code, runs out-of-band, and never blocks the agent.

THE EVIDENCE LADDER

E0 — signed

The receipt exists, signed by a known key.

E1 — in the log

Included in a checkpointed, append-only log.

E2 — anchored

Checkpoint anchored on-chain + independently timestamped.

E3 — co-signed

Counterparty signs the same record: bilateral proof.

Running now, checkable today: the live demo shows 8,902 of 9,229 receipts at E2, anchored hourly (124 anchors and counting, on the Base Sepolia staging network today, mainnet next) · a cold outsider verified a bundle, unaided, in ~2 minutes · flip one byte and the same bundle returns NOT VERIFIED.

THE DATES

JAN 2025

DORA, in force. EU financial firms must keep ICT logs "protected against tampering, deletion and unauthorised access" (RTS 2024/1774, Art. 12(2)(d)). An AI agent is ICT.

DEC 2027

EU AI Act. High-risk AI must automatically record events over its lifetime, logs retained (Art. 12, 19, 26(6)).

2028

The first audits read the 2026–27 records. Evidence cannot be created after the fact.

SEE IT LIVE

The anchored live demo, a verifiable sample Evidence Report, the in-browser verifier, and the AI-Act/DORA mapping — all at swarm.ai

WHO SHOULD TALK TO US

Insurers preparing to underwrite agent risk. Audit firms building agent-assurance practices. Companies putting agents into regulated work. One question tells you if this concerns you: **when your auditor asks how you know your AI logs weren't edited after the fact — what is today's answer?**