ARTEFACT 2 OF 3 · ISACA EUROPE CONFERENCE

The Trust Fabric Blueprint

From Audit Trails to Trust Fabric. Traceability Across Humans and AI Agents

Maman Ibrahim · DiamondSoul · Ginkgo Resilience

Five anchors. Each separately buildable, separately valuable and separately auditable — you do not need all five to start, but you need all five to finish. Underneath them runs one correlation ID. Break the thread at any single anchor and you do not have four fifths of a story; you have no story, because the value was never in the anchors.

1 WHO

Identity

Humans and agents both carry verifiable identity, and every agent identity names a human principal.

STANDARD THAT EXISTS

SPIFFE/SPIRE, CNCF graduated Sept 2022 — workload identity. MCP rev. 2026-07-28: where authorization is implemented, OAuth 2.1 resource server and audience validation are mandatory.

WHERE YOU MUST DESIGN

No ratified agent identity standard, and no SPIFFE profile for AI agents. Specify four fields yourself: agent identity, human principal, delegation depth, revocation point.

2 WHEN

Time

Clocks synchronised to a known source, and the time source itself recorded alongside the timestamp.

STANDARD THAT EXISTS

Delegated Reg. (EU) 2025/1155, in force 2 Mar 2026, grades accuracy by actor speed: 100 µs for HFT, 1 ms for algorithmic trading, 1 s for everything else — plus a documented, annually reviewed system of traceability to UTC (Art 16).

WHERE YOU MUST DESIGN

No AI Act, DORA or NIS2 instrument specifies clock accuracy at all. Set your own tier, then measure and review it.

3 WHAT

Events

One event model across tools and teams, threaded by a single correlation ID that is carried, not recomputed.

STANDARD THAT EXISTS

W3C Trace Context Level 1 — `traceparent`, 16-byte trace-id, a Recommendation since 2021 and probably already in your stack.

WHERE YOU MUST DESIGN

OTel GenAI conventions are status Development. Tool arguments and results are Opt-In — switch them on deliberately.

4 WHY

Decisions

Intent and approval captured as structured records, in the course of deciding, never written up afterwards.

STANDARD THAT EXISTS

For remote biometric ID only, AI Act Art 14(5) requires two-person verification and Art 12(3)(d) requires those verifiers be logged. Narrow in scope; the principle travels.

WHERE YOU MUST DESIGN

The decision record itself. Eight fields, one screen — and the trigger's last three parts: receiver, window, authority.

5 SO WHAT

Outcomes

What changed, what impact occurred, what rollback happened — and whether the rollback was verified.

STANDARD THAT EXISTS

SLSA's ladder transfers cleanly: L1 provenance exists · L2 forging requires an attack · L3 forging requires a vulnerability beyond most adversaries.

WHERE YOU MUST DESIGN

Blast-radius measurement for "wrong at scale, quietly" — the failure where nothing broke and everything was wrong.

One correlation ID threads all five. Break the thread at any anchor and the story stops being a story.

Where the evidence layer now stands

- **RFC 9943 (SCITT)** and **RFC 9942 (COSE Receipts)**, Standards Track, June 2026 — append-only log, registration policy, signed receipts, non-equivocation. Neither mentions AI once; applying them to agent action logs is available and unclaimed.
- **OWASP ASI08** states the requirement outright: "Record all inter-agent messages, policy decisions, and execution outcomes in tamper-evident, time-stamped logs bound to cryptographic agent identities."
- **NIST SP 800-92**, the only final version of NIST's dedicated log-management guide, dates from 2006. Rev 1 has been in draft since October 2023.

Three questions that need no budget

- Are tool arguments and results being captured, or are we running the Opt-In default?
- Does our agent identity name a human principal, and can we revoke it without going through the vendor?
- What is our measured divergence from UTC at the point each log applies its timestamp — *measured*, not assumed?

And the one that decides whether any of it survives: if capturing the field takes longer than making the change, teams will route around it — and a control people route around is worse than no control, because it also reports green.

Reg. (EU) 2024/1689 as amended by Reg. (EU) 2026/1744 (Digital Omnibus on AI, in force 27 July 2026 — Chapter III Sections 1–3 deferred to 2 Dec 2027 for Annex III high-risk and 2 Aug 2028 for Annex I) · Commission Delegated Reg. (EU) 2025/1155 of 12 June 2025, Arts 11–16 and Annex IV, applying from 2 March 2026 and repealing Delegated Reg. (EU) 2017/574 (MiFID II RTS 25) from that date; scope per Art 22c of Reg. (EU) No 600/2014 · Dir. (EU) 2022/2555 Art 23(4) — 24h early warning and 72h notification from becoming aware, final report one month after the notification · Reg. (EU) 2022/2554 Arts 17, 19 and Delegated Reg. (EU) 2025/301 Art 5(1) · ISO/IEC 42001:2023 A.6.2.8 · NIST AI 100-1 and AI 600-1 · W3C Trace Context Level 1 · OWASP Top 10 for Agentic Applications 2026 · RFC 9943, RFC 9942, RFC 8915, RFC 3161. Verified 29 August 2026.

Decision Infrastructure™ | Navigate complexity. Decide with confidence. | thedecisionlayer.online. | maman@thedecisionlayer.online

Maman Ibrahim – ICF Accredited Coach – Fellow ISRM – Fellow IoCR – MIET – ChCSP – CISSP – CISA – CRISC – CDPSE

GINKGO RESILIENCE LTD - 3RD FLOOR, 86-90 PAUL STREET - LONDON - ENGLAND - EC2A 4NE