DiamondSoul
THE DECISION LAYER
Navigate complexity. Decide with confidence.

ARTEFACT 1 OF 3 · ISACA EUROPE CONFERENCE

Minimum Viable Traceability

The seven fields that turn a change log into evidence — plus the four an AI agent must add. From ***"From Audit Trails to Trust Fabric. Traceability Across Humans and AI Agents,"*** Maman Ibrahim.

You do not need perfection to gain value. If a change to a production system cannot show these seven fields, it should not have been made. This is a bar, not a maturity model — and every field belongs in a system you already run. The design constraint that decides whether it survives contact with a delivery team is simple: **if capturing the field takes longer than making the change, people will route around it** — and a control people route around is worse than none, because it also reports green.

The seven fields

	FIELD	WHAT IT MUST CONTAIN	WHERE IT ALREADY LIVES
1	**Actor identity and role**	Who acted, and in what capacity. Human or agent, cryptographic where the system allows it, never a shared account. An agent identity that does not name a human principal is anonymity with extra steps. ***Fails when:*** *a service account, a vendor login or a shared mailbox is the actor of record.*	IAM · agent runtime · pipeline service identity
2	**Trigger event reference**	A pointer to the thing that caused this — the alert, the ticket, the schedule, the upstream failure. Causality is the property auditors most often ask for and logs least often carry. ***Fails when:*** *the reason is "as agreed" or "per the plan", with nothing to follow.*	Monitoring · ITSM · scheduler
3	**Change intent**	One sentence, in the actor's own words, written *before* the change. Not a category from a dropdown — a dropdown captures classification; a sentence captures reasoning, and reasoning is what does not survive a resignation. ***Fails when:*** *intent is inferred later from the diff, or supplied by someone who was not the actor.*	Change record — a mandatory free-text field
4	**Approval reference and scope limits**	Who allowed it, and the boundary that came with the permission. The limit matters as much as the approval: an approval without a stated scope cannot be exceeded, which means it cannot be breached, which means it cannot be tested. ***Fails when:*** *the approval is recorded and the limit is not, or they live in different systems.*	Change record · CAB minutes · agent grant
5	**Pre-check and post-check results**	The state you asserted before, and the state you observed after. Most records have one or the other. Having both is what makes a change log evidence, because the gap between them is where every silent failure lives. ***Fails when:*** *the post-check is "no issues reported", which is an absence, not an observation.*	CI/CD pipeline · deployment verification
6	**Rollback plan and rollback outcome**	Written before. Recorded after. Both, or neither counts. A rollback plan that was never executed and never tested is a paragraph, and paragraphs do not restore systems. ***Fails when:*** *the plan says "restore from backup" with no restore ever having been timed.*	Change record · runbook · pipeline job
7	**Correlation ID**	The thread that ties all six of the above into one story, carried through every system the action touches — not recomputed at each hop by a clever analyst at three in the morning. W3C Trace Context `traceparent` gives you a 16-byte trace-id and is already flowing through most modern stacks. ***Fails when:*** *each system has its own identifier and the join is manual.*	Observability stack · propagate the existing `traceparent`

The two fields people leave out, and why they are the two that matter

Field 3, intent. Every other field can be reconstructed from systems. Intent exists only in a head, and heads resign. A name on a sign-off sheet dated June for a decision made in February is not acceptance — it is retrospective decoration.

Field 6, rollback outcome. The plan is easy to demand and easy to write. The outcome is the only part that proves the plan was real, and it is the field that turns "we could have reverted" into "we did, in eleven minutes".

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

THE AGENT EXTENSION

Four more fields, for anything that acts on its own

An agentic system is a colleague with no employment contract. The seven fields still apply — an agent has an identity, a trigger, an intent, an approval, checks, a rollback and a correlation ID like anyone else. These four are what it adds, and **no observability standard requires any of them today.** Not OpenTelemetry, not ISO/IEC 42001, not the EU AI Act. They are yours to specify, which means they will only exist if you write them into a requirement.

	FIELD	WHAT IT MUST CONTAIN	NEAREST EXISTING HOOK
8	**Why it chose this action**	The alternative considered and not taken. Without a rejected alternative you have a description, not a decision — as true of an agent as of a committee. *This is the field that most often would have prevented the incident, and the one most often absent.*	None. Emit as a span event on the decision step.
9	**What context it used**	The retrieved documents, tool outputs and memory actually present *in that run*. Not the prompt template — the template is documentation; the inputs are evidence. *Retain references and hashes rather than payloads where the content is sensitive.*	`gen_ai.tool.call.arguments` and `gen_ai.tool.call.result` — both Opt-In, see below
10	**Uncertainty, and what it escalated**	Including the escalation it considered and did *not* raise. The near-miss is the most valuable record in an agent estate and almost nobody keeps it, because logging a non-event feels pointless. It is not a non-event — it is your only early-warning surface. *Record the threshold that was not crossed, not only the ones that were.*	`gen_ai.evaluation.score.value` · custom span attribute
11	**The bounded limit it operated inside**	The grant, stated in the trace at execution time — not in a policy document elsewhere. **A limit that is not in the runtime is not a limit. It is a wish.** *Test the boundary by effect, not by name: an agent forbidden to amend a record can still call a tool that amends it.*	Agent grant · policy decision point · gateway

Check this before you specify anything else

In the OpenTelemetry GenAI semantic conventions, `gen_ai.tool.call.arguments` and `gen_ai.tool.call.result` carry the requirement level **Opt-In** — the weakest level the specification has. Default instrumentation therefore records that a tool ran, not *what it did*. A platform can be fully compliant with the convention and useless in an investigation. Turning them on is a configuration change; note that it also makes the trace store a sensitive data store, so decide retention and access at the same time.

Conventions status as at August 2026: Development, not stable. `gen_ai.system` has been renamed `gen_ai.provider.name`, and the conventions have moved to the `open-telemetry/semantic-conventions-genai` repository — any source citing the old name or the old registry page is out of date.

The three moves that fit inside thirty days

Week 1 — run one cold read. Take your last material change or incident decision. Hand the record to someone capable who was not involved. No briefing, twenty minutes, then have them tell it back. Every question they cannot answer is your record's next field.

Week 2 — measure one clock. Pick two systems that both log the same workflow. Ask for the measured divergence from UTC at the point each applies its timestamp. "We're on NTP" is not an answer — it is the finding.

Weeks 3–4 — add three fields. Change intent, approval scope limit, rollback outcome, into the change record you already have. A form change and a policy line. No platform, no vendor, no business case.

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