

Produce records that can stand up to audit, incident response, and challenge

Basic logs may be enough for troubleshooting. They are rarely enough for governance. In agentic workflows, the real test comes later: can a security team, auditor, or regulator reconstruct what happened, defend the integrity of the record, and show who approved what under which policy?

This checklist helps teams evaluate whether an approach produces audit records that are complete, integrity-protected, and usable under scrutiny.

Use this checklist when:

- Deciding whether an agent workflow can pass security review
- Defining evidence requirements for production approval
- Preparing for regulated review, incident response, or internal audit

Checklist (evidence-grade auditability)

- | | |
|--|--|
| <ul style="list-style-type: none"><input type="checkbox"/> Evidence package export: Can you export a package that reconstructs activity end-to-end (events, policy decisions, and relevant session context)?<input type="checkbox"/> Who/what/when/where/why: Do records include actor identity (human/agent), target system, action, result, policy evaluated, and reason codes?<input type="checkbox"/> Trustworthy time: Are timestamps reliable and time-synchronized across components so sequences of actions are defensible?<input type="checkbox"/> Integrity protections: Are records protected from silent modification or deletion (cryptographic sealing, immutability/WORM where warranted, retention locks)?<input type="checkbox"/> Chain-of-custody: Is there a record of evidence handling (access, export, sharing), and can you prove integrity? | <ul style="list-style-type: none"><input type="checkbox"/> Admin activity logging: Are administrative actions on logging, retention, and evidence systems recorded and reviewable?<input type="checkbox"/> Separation of duties: Can you separate operators from evidence administrators and auditors, with least-privilege evidence access?<input type="checkbox"/> Retention controls: Can you configure retention by workflow risk and enforce holds when required?<input type="checkbox"/> Privacy controls: Can you minimize raw capture by default where possible and restrict sensitive evidence access?<input type="checkbox"/> IR readiness: Can incident response use the evidence quickly to answer what happened, what data was exposed, and what actions occurred? |
|--|--|