

Compliance evidence summary: AI coding agents, 2026-10-01 to 2026-12-31

Generated 2026-10-13T00:45:24.598Z by SourceVault 1.58.3 from this install's audit chain, commit history and attestation ledger.

Coverage: audit log on, Sentinel on, provenance on; 1 month(s) of records, chain intact. Policy in force: 3 rule(s), default allow; 2 active agent token(s).

This summary is evidence for your own control assessment. SourceVault holds no SOC 2 report or ISO 27001 certificate; the references name the criterion each line of evidence speaks to.

CONTROL AND WHAT A REVIEWER ASKS	WHAT THIS INSTALL'S RECORDS SHOW	STATUS
Agent access to source is restricted by policy Can an AI agent read any file in the repository, or is access bounded and are refusals recorded? SOC 2 TSC CC6.1 · ISO 27001:2022 A.5.15 · ISO 27001:2022 A.8.3	- Sentinel policy enforcement is on. The policy in force holds 3 rules with default "allow" (current state). 5 reads were refused in the period, across 4 distinct paths; each refusal is a chained record naming the agent, the path and the rule. Verify: Read a path the policy denies through the dashboard or MCP and observe policy_denied; the refusal appears in the audit export as a policy_decision record.	Evidenced
Every agent acts under a named, scoped identity Who read what? Can access be attributed to a specific agent and revoked? SOC 2 TSC CC6.1 · SOC 2 TSC CC6.2 · ISO 27001:2022 A.5.16 · ISO 27001:2022 A.5.18	2 agent tokens are active (current state); 3 mints and 1 revocation were recorded in the period. - Activity was attributed to 3 named agents: claude-code (88), cursor-dana (24), ci-reviewer (21). 13 data-path events carried no identity; those surfaces should be given tokens. Verify: sourcevault agent list shows the registry; every chained record carries a subject field; a revoked token makes the MCP server refuse to start.	Partial
Secrets are redacted before content reaches an agent Could a credential in the codebase reach an AI model through the agent? SOC 2 TSC CC6.7 · ISO 27001:2022 A.8.12 · ISO 27001:2022 A.8.24	- The DLP pass scanned every file read, search result and answer. 5 secrets were scrubbed in 4 redaction events (stripe-secret-key 2, aws-access-key-id 1, dlp 1, generic-api-key 1). - Redaction happens inside the engine, before the content is returned to the client and therefore before a cloud-backed client could forward it. Verify: Read a file containing a test credential through MCP and observe the [[redacted]] marker; the scrub is a policy_decision record with the DLP rule name.	Evidenced
Agent activity is logged in a tamper-evident record Is there a complete, protected log of what AI read, and can it be exported to our SIEM? SOC 2 TSC CC7.2 · SOC 2 TSC CC7.3 · ISO 27001:2022 A.8.15 · ISO 27001:2022 A.8.16 · ISO 27001:2022 A.5.33	- The audit log is on. 153 events were recorded in the period across 1 month: 58 searches, 9 asks, 68 file reads. - Every month's hash chain verifies, each anchored by the signed head. - The chain exports as JSONL with its hash fields, CEF or syslog; a month that fails verification makes the export exit non-zero. Verify: sourcevault audit verify, or GET /api/dashboard/audit/verify; export with sourcevault audit export and recompute the chain on your side.	Evidenced
AI-authored changes to security-relevant code are identified and attested Which commits did AI write, which of those touched authentication, secrets, payments or infrastructure, and is that record signed? SOC 2 TSC CC8.1 · ISO 27001:2022 A.8.32 · ISO 27001:2022 A.8.25	12 commits across 2 repositories in the period; 7 disclosed AI authorship; 6 of those touched security-relevant code (authentication 2, ci_supply_chain 1, dependencies 1, network_egress 1, payments 1). 1 signed attestation was recorded; the attestation ledger verifies. - AI authorship is a disclosure signal (commit trailers, footer markers, bot identities), not a forensic finding. Verify: GET /api/dashboard/repos/<repo>/provenance reproduces the classification from git log; attestations are DSSE envelopes verifiable against attestation_key.pub.	Evidenced
Changes to the security configuration are themselves recorded Could someone switch the controls off without a trace? SOC 2 TSC CC6.8 · SOC 2 TSC CC8.1 · ISO 27001:2022 A.8.9	1 security settings change was recorded in the period and 0 were refused by the license gate; turning the audit log off is itself logged before it stops. 0 license installs, 0 dashboard token rotations, 1 authentication failure in the period. Verify: settings_change, license_install and token_rotate records in the audit export; the chain verdict covers them like every other record.	Evidenced

Sources: the AI activity report for the same period (sourcevault report ai-activity) states each section's coverage and how to reproduce every figure. Current-state values were read when this summary was generated. Status: Evidenced means the control was on and the records show it acting; Partial means it was on but the period holds little or incomplete evidence; Not evidenced means the control was off or its record failed verification.