@demotoprod · CHECKLIST

The agent audit-trail checklist

A flight recorder keeps at least the last 25 hours of flight data and must survive 3,400 g of impact, 1,100 °C of fire for an hour and 30 days at a depth of 20,000 feet. An AI agent's usual record is a transcript and a summary the model wrote about itself: the only witness, writing its own report. The EU AI Act makes providers of high-risk systems keep the logs under their control for at least six months; India's national cyber agency recommended in October 2026 logs kept at least 180 days and a tamper-resistant log of agent activity. Five checks that turn a transcript into evidence.

1 Write the log at the gateway, not in the model

The model that acted is the only witness, and it writes its own summary. A record is evidence only when the instrument that writes it is not the actor. India's national cyber agency asks that model-initiated requests be logged as completely as feasible, including destination, port, protocol and result, and that every action affecting production systems and data be in a tamper-resistant log.

- Every tool call passes through one layer you control (the tool layer, the API proxy or the approval step); that layer opens the line before the call and completes it after.
- Per call: run id, step, timestamp, the tool and its arguments, the result or the error, model id and version, prompt version, and who approved.
- Append-only and hash-chained: each line carries the hash of the line before it, so a changed or missing line shows.

```
{"run":"r-7f2a","step":14,"ts":"2026-10-19T21:07:12Z",
 "tool":"delete_rows","args":{"table":"orders","where":"stale"},
 "result":"ok","rows":147,"model":"v2026-10","prompt_v":"22",
 "approver":"ops-lead","prev":"sha256:9c1e…","mode":"append-only"}
```

2 Keep it six months at least, where the agent cannot write

The AI Act makes providers of high-risk systems keep the logs under their control for a period suited to the system's purpose, and at least six months unless other law says otherwise (Article 19). India's national cyber agency, in an advisory, asks for at least 180 days. A log the agent can reach is a log the agent can tidy.

- Retention is a number in a config, set to the longest floor that applies to you and never below six months; not a disk-full default.
- The store is write-once for the gateway and read-only for everyone else; the agent's credentials cannot list, read or delete it (a separate account, separate keys, object lock).
- Test the lock, not the policy: attempt the delete with the agent's credentials, and keep the refusal in the log.

```
retention_days: 184      # six months is at most 184 days; India's cyber agency: 180
agent_role: {log: deny}  # the actor never holds a key to its own record
```

3 Replay from the log alone, every month

Investigators download the recorder's data and analyse it with specialised software; the crew's account is not the record. India's national cyber agency warns that after-the-fact transcript review alone can miss activity that a model's reasoning makes appear legitimate. If a run cannot be rebuilt from the log with the model switched off, the log is a diary.

- Pick one real run from the last month; from its log lines alone, rebuild the sequence of tool calls, their inputs and results, and compare with what the systems show.
- A gap, an unlogged call or a line nobody can explain is an incident, filed and fixed before the next drill.
- Put the drill on the calendar with a named owner; time it, and track the time as a release metric.

4 Record the approval with the action

An action without its approver is half a record. When the question is 'who let it do that', the answer has to be in the same line, not in a chat thread. The AI Act already asks this of remote biometric identification: its logs must identify the people who verified each result (Article 12).

- Every action above your risk line (writes, deletes, payments, messages sent) carries the approver's id, the policy version that allowed it and the time of approval.
- Automatic approvals are logged as such, with the rule that granted them; silence is never approval.

5 Count the gap every day

The failure that hides is the call that never reached the log. In the cases India's national cyber agency reported in October 2026, none of the affected organizations detected the agents' access; each was disclosed by the model provider after its own review. One number, read daily, finds the gap before an auditor does.

- Compare the tool layer's call count with the log's line count for the same day; any difference is an incident.
- Alert when the log goes quiet while the agent is running, and when a line arrives with a broken hash chain.

```
calls=$(tool-layer count --day today)
lines=$(wc -l < agent-log/today.jsonl)
[ "$calls" -eq "$lines" ] || alert "unlogged actions: $((calls-lines))"
```

Nitish Gautam · indica-tech.com · hello@indica-tech.com

INDICA TECH · @DEMOTOPROD

Nitish Gautam

Founder-led AI engineering. Demo to production, under NDA, in regulated industries, on fixed scope.

Paste into your project

A LOG POLICY TO ADAPT AND PASTE

```
log_writer: gateway       # the tool layer writes each line; never the model
fields: [run, step, ts, tool, args, result, model, prompt_v, approver, prev]
chain: sha256             # each line carries the hash of the line before it
store: write-once         # object lock; the agent's credentials cannot reach it
agent_role: {log: deny}
retention_days: 184       # six months is at most 184 days; India's cyber agency: 180
replay_drill: monthly     # one real run, rebuilt from the log alone
gap_check: daily          # tool-layer calls == log lines, or it is an incident
```

Where this list comes from

Ten-plus years taking products from zero to one across seven sectors, including defence, fintech, medical and telecom, with AI shipped into regulated and air-gapped environments. The @demotoprod channel takes one production AI failure apart each day. These checklists are the same work written down: every item is a dated failure or a documented default, with the control that stops it and the primary source.

If this list found a gap

1 **Your demo works and production is stuck.**

Production Readiness Audit. Fixed price, two weeks. A five-dimension scorecard, cost and failure analysis, and a 90-day execution roadmap.

2 **You need senior AI judgement without a full-time hire.**

Fractional AI Lead. Monthly: 90-day minimum, then 30 days' notice either side. Weekly 1:1 with the founder or CTO, architecture and code review, technical hiring support, investor-ready due diligence.

3 **An AI feature has to work, by a date.**

AI Build Sprint. Fixed price, four to eight weeks. Scope signed off up front, weekly demos, CI/CD and observability included, 30-day production warranty.

4 **A regulator, an auditor or a customer's security team is asking.**

AI Governance & Compliance. Fixed price, four weeks. ISO 27001 / SOC 2-aligned controls, a model risk framework, human-in-the-loop and audit logging, regulator-ready documentation.

BOOK A CALL
indica-tech.com/#contact

EMAIL
hello@indica-tech.com

WEBSITE
indica-tech.com

ALL CHECKLISTS
indica-tech.com/checklists/

YOUTUBE
youtube.com/@demotoprod

LINKEDIN
linkedin.com/in/startup-founderx

INSTAGRAM
instagram.com/demotoprod

TIKTOK
tiktok.com/@demotoprod

Free to share. The live list, with every source link: indica-tech.com/checklists/#agent-audit-trail