

Decisions · Platform

ADR-0009: Production is verified continuously, and logs never carry what people search

🕒 2 min read · Updated 25 September 2026

How production is verified continuously by a data canary and request ids, and why logs never carry what people search.

<https://docs.gurbanisoul.com/adr/0009-production-verification/>

Status: accepted (2026-09-24)

Context

Deploy gates prove a release when it ships. Afterwards a stale CDN copy, a wrong database behind the API or a partial rollback could serve the wrong text with nothing noticing. And tracing a request from Vercel to the API was guesswork, while the privacy policy promises no tracking is added to server logs.

Decision

- **Data canary** (tools/data_canary.py, workflow data-canary, every 6 hours): 500 random lines and 12 random Angs plus Angs 1, 712 and 1430, fetched from the API origin and through the public site's CDN, compared byte for byte with the pinned database; the golden contract replays against production. A difference opens one issue; the printed seed replays it.
- **Request ids**: every response carries X-Request-Id (a valid incoming id or Vercel's x-vercel-id, else a random one; anything unsafe is replaced), and the access log records it.
- **Privacy-safe logs**: one JSON line per request with method, path, status, duration and the id — never the query string, IP or user agent. A startup line records version, commit, dataset and modules.

Consequences

Wrong or stale scripture in production is detected within hours, not by a reader. A request can be followed across Vercel and Render without logging what anyone searched for.