FREE RESOURCE · ROGER GV — AI PRODUCT BUILDER

AI App Launch Audit Checklist

The 25 checks I run before any AI-generated or inherited app goes to production. Most AI-built apps pass functional tests and still ship critical vulnerabilities — this list is how you catch them first.

How to use it: go through each check honestly. **"Implemented" is not "tested", and "tested" is not "production-ready".** If you can't verify a check with evidence (a test, a log, a rule you can point to), it's not done. Three or more open items in sections 1–3 = do not launch yet.

01 Authentication & sessions

- [ ] **Every route that shows user data requires authentication.** Open a private URL in an incognito window — it must redirect, not render.
- [ ] **Sessions are isolated per user.** Log in as two users in two browsers; user A must never see user B's data (the #1 vibe-coding bug).
- [ ] **Session tokens are signed and HTTP-only.** No auth state living only in localStorage or client-side variables.
- [ ] **Password reset / magic links expire and are single-use.**
- [ ] **CSRF protection exists on every state-changing form or endpoint.**

02 Data security (the RLS section)

- [ ] **Row Level Security (Supabase/Postgres) or security rules (Firestore) are enabled on every table/collection.** Default-open databases are the most common critical finding in AI-generated apps.
- [ ] **Rules are tested with a non-owner account — not just written.** Query another user's records directly via the API; it must fail.
- [ ] **The client never receives fields it doesn't need** (emails, tokens, internal flags in API responses).
- [ ] **File storage (buckets) has per-user access rules.** Try fetching another user's file URL without their session.
- [ ] **Personal data has a deletion path.** If a user asks you to delete their account, can you actually do it?

03 State authority & payments

- [ ] **The server — never the client — is the authority on money, quantity and status.** Prices, balances and booking states validated server-side; the client only displays.
- [ ] **Payment state changes only on webhook confirmation** (Stripe PaymentIntent succeeded), not on the client's "success" screen.
- [ ] **Webhooks verify signatures** (Stripe signing secret, HMAC on custom hooks).
- [ ] **Critical operations are idempotent.** A retried webhook or a double-click must not create two orders or two charges.
- [ ] **Race conditions on the critical flow are handled** (two users booking the last slot at the same time).

04 Secrets & configuration

- [ ] **No API keys or secrets in the client bundle or repo history.** Search the built JS for "sk_", "key", "secret". AI tools paste keys into frontend code constantly.
- [ ] **Dev and prod use separate keys, databases and projects.**
- [ ] **Debug endpoints, seed routes and test users are removed or gated.**
- [ ] **Error messages don't leak internals** (stack traces, SQL, file paths) to the user.
- [ ] **Dependencies audited:** no known-critical CVEs, no abandoned packages holding auth or payments.

05 AI-specific checks (if your app calls an LLM)

- [ ] **User input is treated as untrusted in prompts** — prompt injection can't make the model reveal other users' data or trigger privileged actions.
- [ ] **Consequential actions (send, pay, publish, delete) require human approval** — the model proposes, a person disposes.
- [ ] **Model outputs are validated before use** (schema-checked, never executed or inserted raw).

06 Release readiness

- [ ] **The critical flow has automated tests** — at minimum the money path and the auth path.
- [ ] **You can roll back.** A tagged previous version, database backup, and a tested restore procedure.
- [ ] **Logs let you answer "what happened?"** for any user action in the critical flow.
- [ ] **Someone is the explicit owner of the go/no-go decision** — and it's written down, with the evidence that justified it.

Scoring your audit

0–2 open items (none in sections 1–3): you're in launch territory — fix the rest in week one.
3–5 open items or any open item in sections 1–3: launch to a closed beta only, with real monitoring.
6+ open items: no-go. Not because the app "doesn't work" — because you can't yet prove it does. That proof is the whole job.