The AI App Production Checklist

47 checks I run on every Lovable / Bolt / Base44 / Replit build before it faces real users — by Mohit Sengar, fractional CTO (mohitsengar.org)

1 · Authentication & Sessions

- ■ 1. Sign-up, login, logout, and password reset all work for a SECOND user account — not just yours
- ■ 2. Sessions expire; refresh tokens rotate; logout actually invalidates the session
- ■ 3. OAuth redirect URLs locked to your domains (no wildcard localhost in production)
- ■ 4. Email verification required before sensitive actions
- ■ 5. Rate limiting on login and password-reset endpoints (credential stuffing)
- ■ 6. Admin routes gated by role checks on the SERVER, not hidden buttons in the UI

2 · Supabase / Database Security

- ■ 7. Row-Level Security enabled on EVERY table — not just the obvious ones
- ■ 8. RLS policies tested as another user: can user B read user A's rows?
- ■ 9. Service-role key never shipped to the browser bundle (search your JS for it)
- ■ 10. Database has automated backups with a tested restore path
- ■ 11. Foreign keys and NOT NULL constraints exist (AI often skips them)
- ■ 12. No orphan rows after deletes — cascade or clean up explicitly
- ■ 13. Indexes on every column used in WHERE / ORDER BY at scale

3 · Secrets & Configuration

- ■ 14. No API keys hard-coded in the repo or client bundle — check git history too
- ■ 15. Separate keys for dev / staging / production
- ■ 16. Environment variables documented; a new engineer can boot the app from README alone
- ■ 17. .env files in .gitignore — and rotated if ever committed

4 · Payments (Stripe & friends)

- ■ 18. Webhook signatures verified — not just parsed
- ■ 19. Idempotency keys on charge creation (double-click ≠ double-charge)
- ■ 20. Webhook retries handled: duplicate events do not duplicate records
- ■ 21. Refund and dispute flows exist — even if manual
- ■ 22. Test-mode keys cannot reach production and vice versa
- ■ 23. Price/plan amounts come from the server, never trusted from the client

5 · APIs, Webhooks & Integrations

- ■ 24. Every external call has a timeout and a retry-or-fail plan
- ■ 25. Third-party failures degrade gracefully (Twilio down ≠ app down)
- ■ 26. Inbound webhooks authenticated (signature or secret) and replay-protected
- ■ 27. LLM calls have max-token caps, cost tracking, and prompt-injection guards
- ■ 28. Rate limits on your own public endpoints

6 · Data Integrity & Edge Cases

- ■ 29. Account deletion works and cascades correctly (GDPR baseline)
- ■ 30. Empty states, long strings, emoji, and RTL text don't break the UI

- ■ 31. Concurrent edits don't silently lose data
- ■ 32. File uploads validated: type, size, and storage rules (who can read them?)
- ■ 33. Timezone handling explicit — dates stored UTC, rendered local

7 · Observability & Operations

- ■ 34. Error tracking wired (Sentry or similar) with alerts to a human
- ■ 35. Structured server logs you can actually search
- ■ 36. Uptime monitoring with a public or internal status check
- ■ 37. A staging environment that mirrors production
- ■ 38. Deploys are repeatable (CI/CD) — not drag-and-drop from a laptop
- ■ 39. Rollback path documented and tested once

8 · Performance & Scale Sanity

- ■ 40. Largest page loads under ~3s on a mid-range phone over 4G
- ■ 41. N+1 query hot spots checked on list views
- ■ 42. Images compressed and sized (no 4MB hero PNGs)
- ■ 43. Caching strategy exists for expensive reads

9 · Handover & Legal Hygiene

- ■ 44. Code lives in YOUR repo; infra in YOUR cloud accounts
- ■ 45. A written runbook: how to deploy, restore, rotate keys, and page someone
- ■ 46. Privacy policy and terms exist and match what the app actually does
- ■ 47. Analytics/tracking disclosed; cookie consent where required

Scoring: 40+ checked — ship it. 30–39 — fix the gaps first. Under 30 — book a build review before launch:
calendly.com/mohitsengarr · mohit@sengarconsultancy.org · © Mohit Sengar / Sengar Consultancy 2026