Security & Trust Overview

CheckSmith.ai™ — maintained by the CheckSmith team. Last updated July 22, 2026.

Shared Responsibility Model

Security is a partnership across three layers:

- **Platform (Lovable Cloud):** managed hosting, database, authentication, and network infrastructure — including TLS termination, encryption at rest, and platform patching.
- **Application (CheckSmith.ai):** row-level security policies, role-based access, workspace isolation, audit logging, and secure server-side business logic.
- **Customer:** account credentials, workspace PIN, role assignments, team membership, and the data you choose to import or share.

Access & Authentication

- Email/password and Google sign-in via managed authentication; no anonymous access.
- Workspace-level PIN gate with server-hashed secrets — PIN hashes and salts are never sent to the browser (excluded from real-time broadcasts as of the latest release).
- Role-based access control (admin / member) enforced by a security-definer `has_role` function; roles stored in a dedicated table, never on user profiles.
- Session-expired detection and automatic sign-out on the admin console.

Data Isolation & Row-Level Security

- Every tenant table has row-level security enabled with policies scoped to workspace membership.
- Real-time subscriptions are column-scoped — sensitive columns (e.g. client PIN hashes) are excluded from broadcast publications.
- Server functions authorize with the caller's session before privileged operations; the service-role client is used only inside verified server-only modules.

Encryption & Transport

- All traffic served over HTTPS/TLS; HSTS enabled on the published domain.
- Data at rest is encrypted by the managed cloud platform.
- Secrets (API keys, webhook signing keys) stored in the platform secret manager, never in source or client bundles.

Auditing & Monitoring

- Admin actions, workspace changes, signup events, client-PIN attempts, and posting failures are written to append-only audit tables.
- Visitor and analytics events are logged with IP/user-agent for abuse detection; retention policies archive older records.
- Automated dependency and configuration scans run against the codebase; recent findings (real-time PII exposure, delete-policy scoping) are remediated and tracked.

AI Privacy

- AI features (invoice extract, categorization, chat) run through the Lovable AI Gateway; prompts and responses are not used to train third-party models.
- Customer content sent to AI providers is scoped to the requesting workspace and never joined across tenants.

Incident Response & Reporting

Suspected vulnerabilities can be reported to **security@checksmith.com**. We acknowledge reports within two business days and coordinate disclosure with the reporter. Verified incidents are triaged, patched, and — where required — communicated to affected customers.

Compliance Status

CheckSmith.ai is not currently certified against SOC 2, ISO 27001, HIPAA, or PCI DSS. We operate on a SOC 2-aligned control baseline (tracked in an internal readiness checklist) and can share our shared-responsibility documentation on request. This page is app-owned editable content, not an independent audit.