Security Program Overview

How probuck.ai protects your documents, data, and identity.

At a glance
probuck.ai is built around a strict Bring-Your-Own-Key model, mandatory two-factor authentication, encryption in transit and at rest, and a fully audited retrieval pipeline. Your documents and AI calls stay isolated to your organisation.

1. Data Privacy & Tenant Isolation

Every customer organisation is logically isolated end-to-end. Documents, embeddings, extracted tables, workflows, chat history, and audit logs are tagged with an organisation ID and protected by row-level security in the database layer.

- Per-organisation isolation enforced by Postgres row-level security on every table.
- Encryption in transit (TLS 1.2+) for all client and server-to-server traffic.
- Encryption at rest (AES-256) for object storage and database volumes.
- Organisation-supplied API keys encrypted with server-side pgcrypto before storage.
- No cross-tenant data access — platform admins are settings-only on other organisations.

2. Access Control & Authentication

Authentication is handled by a managed identity provider with bcrypt-hashed credentials in a reserved schema that the application never reads directly.

- Mandatory TOTP-based two-factor authentication for every user.
- Sessions stored exclusively in sessionStorage and cleared on sign-out.
- 60-second inactivity timeout on the MFA challenge screen.
- Optional IP allow-listing per organisation, enforced on every authenticated request.
- Invite-only organisation membership with strict one-organisation-per-user enforcement.

3. Secure Software Development Lifecycle

All changes to the platform are reviewed, tested, and shipped through an automated pipeline with security gates at each stage.

- Code review required on every change before deployment.
- Automated dependency scanning and vulnerability monitoring.
- Linting, type-checking, and unit tests run on every build.
- Database migrations versioned and reviewed before promotion.
- Edge functions deployed atomically with automatic rollback on failure.

4. Infrastructure & Hosting

probuck.ai runs entirely on managed cloud infrastructure with no self-hosted components, eliminating an entire class of patching and configuration risk.

- Application hosted on Lovable's managed edge platform with global CDN.
- Backend services (database, auth, storage, edge functions) on Lovable Cloud.
- Daily encrypted database backups with point-in-time recovery.
- Three-tier disaster recovery with automated pg_dump and Auth Admin API exports.
- No customer data persisted to local disk on edge workers.

5. Visibility & Auditability

Every retrieval, workflow run, and authentication event is recorded in an immutable audit trail. Customers can inspect exactly how each AI answer was produced.

- 365-day service-role audit log covering authentication and activity events.
- Per-answer diagnostics trail exposing every retrieval and generation step.
- Source citations on every AI answer, traceable back to the originating document and page.
- Token usage tracked per workflow step and per organisation.
- Terms-of-service acceptance recorded with timestamp and IP address.

6. AI Data Governance

AI processing follows a strict Bring-Your-Own-Key model. Your documents are sent only to the AI provider you have authorised, using credentials you control.

- Strict BYOK isolation — no platform fallback keys for organisations.
- Supported providers: Google Gemini and Gemini Enterprise / Vertex AI.
- API keys are organisation-scoped, encrypted at rest, and never exposed to other tenants.
- Document content is not used to train any third-party model.
- Charts and diagrams are excluded from ingestion; only text and structured tables are indexed.

A note on AI accuracy
AI-driven retrieval is powerful but not infallible. probuck.ai is designed for auditability — every answer is traceable to its source so your team can verify the workings. We do not claim AI retrieval is guaranteed accurate.

7. Incident Response

Security incidents are triaged, contained, and communicated to affected organisations on a defined timeline.

- On-call rotation with defined escalation paths.
- Customer notification for incidents materially affecting their data.
- Post-incident review with documented remediation actions.
- Coordinated disclosure process for externally reported vulnerabilities.

For enterprise-specific requirements, data-handling addenda, or further questions, please reach out via the contact form on probuck.ai.