

KeptPDF Technical Privacy Brief

For HIPAA Compliance Review

Author: KeptPDF — Founder

Date: 2026-06-10

Pack version: 1.0

Disclaimer: This document describes KeptPDF's technical architecture as of the version published on the cover date. It is not legal advice.

1. Summary

KeptPDF is a browser-based suite of PDF tools (redact, sign, merge, split, OCR, compress, convert, Bates, ...). All PDF processing happens inside the user's browser. KeptPDF's servers do not receive, transmit, or store PHI in the course of normal tool operation.

2. Architecture

2.1 Client-side processing

- PDF parsing: PDF.js (Mozilla, open source)
- PDF editing: pdf-lib (open source)
- OCR: Tesseract.js (WebAssembly, runs in browser)
- Detection / redaction: KeptPDF's own JavaScript, served as static assets

2.2 Server-side surfaces

- Static asset hosting (HTML, JS, CSS, images, sample files)
- Magic-link authentication (email-only, no password)
- Subscription billing (Stripe)
- Anonymous and signed-in usage counters (counts only, no content)
- HIPAA Compliance Pack email capture (this pack)

3. Data flow

3.1 What stays in the browser

- User-uploaded PDFs and all extracted text
- PII detection results
- Editor state (rectangles, drawing actions, undo/redo)
- Redacted output PDFs and audit certificates
- History (IndexedDB; local-only, never synced)

3.2 What reaches KeptPDF servers

- The user's email address (if signed in)

- The user's IP address (salted-hashed for rate limiting)
- Per-tool usage counts (tool name only — not file content, name, size, or hash)
- Stripe billing metadata (Stripe handles payment data directly)

4. Subprocessors

- **Vercel** (hosting)
 - Sees: page requests, IP addresses, asset bytes served
 - Does NOT see: PDF contents, extracted text
- **Stripe** (billing)
 - Sees: customer + payment data
 - Does NOT see: PHI
- **Resend** (transactional email)
 - Sees: email addresses, message bodies (magic links, this pack confirmation, billing receipts)
 - Does NOT see: PHI

5. Retention

- Auth sessions: 30 days (browser cookie)
- Usage counters: 90 days rolling
- Stripe customer data: retained per Stripe defaults
- Compliance Pack email capture: retained until opt-out

6. Security

- All traffic is TLS 1.2+ (Vercel default)
- Magic-link tokens are single-use, 15-minute TTL
- No password storage
- Source available on request for security audit

7. Limitations

- Browser-based processing depends on the user's device security. If a user's device is compromised, PHI exposure is the user's responsibility, not KeptPDF's.
- KeptPDF's redact auto-detect covers 11 of the 18 HIPAA Safe Harbor identifiers; manual review is required for full Safe Harbor de-identification. The audit certificate documents this coverage explicitly per-document.

Contact

support@keptpdf.com — <https://keptpdf.com>

Signed

KeptPDF — Founder
support@keptpdf.com