

MODERN PAPER, INC. · SECURITY WHITEPAPER

How we protect your data.

A summary of the technical and organizational controls that protect your data when you process PDFs through the PDF Blocks API — written for the security and procurement teams who evaluate us.

VERSION

1.4

DATE

27 July 2026

APPLIES TO

PDF Blocks API

CONTACT

security@pdfblocks.com

01 Overview

This whitepaper covers the **PDF Blocks document-processing API** — the stateless service that receives, transforms, and returns your PDFs, where your data is handled. It runs separately from our account and billing systems, which never receive your documents. Most of what makes PDF Blocks safe to adopt comes down to two structural facts about how the API works:

-
- 01 **We never store your documents.** The API processes each file in memory and returns the result in the same response. Nothing is written to a database or durable storage.
-
- 02 **Our pipeline is deterministic, with no LLMs.** Every action is document manipulation implemented in code, producing the same output for the same input. No language model reads your documents, and nothing you send us is used to train or fine-tune a model.
-

02 Stateless by design

The API that transforms your PDFs has **no database and no durable storage**. Each document is processed in memory and discarded when the response is returned; for each request we retain only metadata (the action, byte counts, an API-key digest, status, and timestamp) — never document contents. It runs in the region you address.

Account management — the dashboard and billing — is a separate system that never receives your documents and is out of scope for this document.

03 Data protection & retention

Documents are received over TLS, processed in memory, and streamed back. Nothing about a document's content is retained. For each request we record only metadata — an API-key digest, the action, byte counts, status, and a timestamp — used for billing, security, and abuse prevention. Because documents are never persisted, there is no document data to export, delete, or produce a certificate of destruction for; the guarantee is structural.

04 AI & machine learning

PDF Blocks is not an AI product. Actions such as merge, split, rotate, watermark, stamp, and password protection are deterministic transformations written in code, and they are the whole of what the API does with your file. No large language model reads your documents, and no document content is sent to a third-party AI service.

A few operations work on what a page looks like rather than on its structure. Reading a barcode is the clearest example: locating the symbol on the page and decoding it may use visual computing. That is image analysis, not a language model, and it runs inside the same isolated environment with no outbound internet access as everything else. Either way the rule is the same: we do not train machine-learning

models on your documents, and we never use their contents for anything beyond the operation you request.

05 Encryption

- **In transit.** All connections use TLS; we enforce a modern TLS 1.3 / 1.2 policy and redirect plain HTTP to HTTPS. Internal service-to-service traffic is also TLS-encrypted.
- **At rest.** The API keeps no durable storage, so there is nothing to encrypt at rest by design. A large upload may be briefly buffered to a temporary file for the duration of a single request and is deleted the moment the request completes. API keys are stored only as one-way digests, never in reversible form.

06 Authenticating to the API

- **API keys** authenticate every request and identify your account for usage metering. Each key is stored only as a SHA-256 digest with a short display prefix, shown once at creation and never recoverable afterward, and can be revoked instantly.
- **A key is a usage credential, not a door to your data.** There is no endpoint that lists, downloads, or replays documents you have already processed, because we never stored them. Someone who steals a key can spend your quota, which is why keys can be revoked instantly, but **cannot reach a single one of your documents**.
- **Least privilege** throughout: each service component runs with a narrowly scoped identity, and internal services are network-segmented so they are reachable only by the components that need them.

07 Infrastructure & network isolation

The API runs on Amazon Web Services and Microsoft Azure, defined as infrastructure-as-code and deployed through a reviewed, automated pipeline.

- **Network isolation.** Document processing runs in isolated private networks with **no outbound internet access**, so it cannot exfiltrate data even in principle.
- **Edge protection.** The API is protected by a web application firewall with rate limiting.
- **Hardened, ephemeral hosts** built from current patched images; instance disks are encrypted and destroyed on retirement.
- **Physical security** of the data centers is inherited from AWS and Microsoft Azure, which maintain SOC 2, ISO 27001, and PCI DSS attestations.

08 Regions & data residency

You choose where documents are processed by addressing a regional endpoint; the request is handled entirely within that jurisdiction. Available regions, each stateless and running the full action catalog:

REGION	ENDPOINT
Global (latency-routed)	api.pdfblocks.com
United States	us.api.pdfblocks.com
United States (HIPAA-aligned)	hipaa.api.pdfblocks.com
European Union	eu.api.pdfblocks.com
United Kingdom	uk.api.pdfblocks.com
Canada	ca.api.pdfblocks.com
Australia	au.api.pdfblocks.com
Japan	jp.api.pdfblocks.com
India	in.api.pdfblocks.com
Brazil	br.api.pdfblocks.com

This choice governs where documents are processed — each request stays entirely within the region you select, and nowhere else.

09 Application security & development

- Every change is version-controlled, peer-reviewed, and merged through continuous integration that runs unit, integration, and cross-service contract tests before shipping.
- Each commit is scanned with static application security testing and dependency-vulnerability auditing; dependency updates are proposed automatically and reviewed.
- Deployments authenticate with short-lived, federated credentials (no long-lived cloud keys) and require an explicit human approval step to reach production.
- Releases are health-checked and roll back automatically on a failed readiness self-test.

10 Vulnerability management & incident response

Vulnerabilities are identified continuously — dependency scanning, static analysis, provider advisories, and external reports — triaged by severity, and remediated on defined timelines, with critical issues addressed within days. We publish a security contact at pdfblocks.com/.well-known/security.txt and welcome coordinated disclosure.

Security incidents are handled under a documented plan with defined severities and response steps. We notify affected customers of confirmed or suspected customer-impacting incidents without undue delay, and — where we act as a processor — notify the customer as controller of a relevant personal-data breach as set out in our Data Processing Addendum.

11 Availability & resilience

- In every region, the API runs with redundant capacity across multiple availability zones, with health checks and automatic replacement of unhealthy capacity.
- Deployments use rolling updates with automatic rollback, so a failed release does not take the service down.
- Because the API is stateless, recovery is a matter of healthy capacity — there is no document data to restore — and a disruption in one region does not affect the others.

12 Subprocessorsors

Document processing runs on our own infrastructure with a small, vetted set of cloud providers:

PROVIDER	PURPOSE	LOCATION
AWS & Microsoft Azure	Cloud hosting and compute	The region you address

The contents of documents you send to the API are not shared with any subprocessor. The full list is maintained at pdfblocks.com/subprocessorsors.

13 How our controls map to common frameworks

PDF Blocks is engineered to align with the criteria that formal frameworks such as SOC 2 and ISO 27001 assess. The table below maps our controls to the SOC 2 Trust Services Criteria; we are glad to walk a reviewer through the evidence behind any of them.

CRITERION	HOW PDF BLOCKS ADDRESSES IT
Security	TLS 1.3; web application firewall and rate limiting; isolated private networks with no outbound access; hashed API keys; least-privilege access; automated vulnerability scanning.
Availability	Redundant capacity across availability zones, health checks, a multi-region API, and rolling deploys with automatic rollback.
Confidentiality	Documents processed in memory and never stored; encryption in transit; strict network segmentation; confidentiality obligations on staff and subprocessors.
Processing Integrity	Deterministic document operations; peer review, CI, and contract tests gating every release; change management with health-checked rollout.
Privacy	No document retention; documents never shared with any subprocessor; no use of customer data to train AI; data residency selectable per region.

14 Requesting documentation

Business customers and prospects under a mutual NDA can request our internal security policies — Information Security, Access Control, Cryptography, Data Protection & Retention, Vulnerability Management, Secure Development, Incident Response, and Business Continuity — and a completed security questionnaire. We also support non-intrusive assessments by arrangement.

To begin, contact **security@pdfblocks.com** or visit **pdfblocks.com/security**.