

Privacy Policy

How we collect, use, share and protect your data — written to be read, not skimmed.

Global Copilot LLC

Effective Date: January 1, 2026 · Last updated: September 22, 2026

See also our [General Terms & Conditions](#) and our [Security & IT Overview](#).

1. Who We Are

Global Copilot LLC ("Global Copilot," "we," "us") operates GlobalCopilot.com (the "Platform"), a market intelligence tool that helps companies research and enter new markets.

Address: 3955 Canyon Road, Lafayette, CA 94549, USA

Contact: hello@globalcopilot.com

For purposes of GDPR and UK GDPR, Global Copilot LLC is the data controller.

2. What We Collect

Account Data

When you sign up, we collect your name, email address, job title, and organization details. If your organization maps its team structure on the Platform, we store that too.

Content You Create

Everything you build on the Platform stays yours. This includes files you upload, research you edit, expert lists you create, and any AI-generated analysis the Platform produces for you.

Technical Data

When you use the Platform, we automatically collect your IP address, browser type (via user-agent header), referring URL, and a session identifier to tie your activity together. We record page views and instrumented interactions (button clicks, form submissions, search queries) and how long you spend on each page; when you are signed in, this activity is associated with your account.

We do **not** collect device fingerprints or hardware identifiers.

Payment Data

Stripe handles all payment processing for subscriptions and one-time purchases via Stripe Checkout. We never see or store your card number.

Data Minimization and Purpose Limitation

Global Copilot adheres to the principles of data minimization and purpose limitation under Article 5(1)(b) and 5(1)(c) of the GDPR. We collect only the personal data that is adequate, relevant, and limited to what is necessary for the specific purposes described in this Policy. Personal data is not processed in a manner incompatible with those purposes.

Where technically feasible, we:

- Restrict internal access to personal data based on role
- Retain personal data only for as long as necessary
- Pseudonymize or aggregate data used for analytics and product improvement

3. How We Use Your Data

We use your data to run the Platform: generating market analyses, processing your files, enabling team collaboration, handling payments, sending you service emails, and keeping things secure.

We also use aggregated usage data to improve the product, and we may send occasional marketing emails (you can opt out anytime).

4. How AI Works on Our Platform

This is the section most people want to read carefully, so we'll be specific.

Market Research AI

We use **Perplexity AI** for market analysis, competitor research, question answering, and content generation.

When you ask the Platform a question or trigger an analysis, your prompt and relevant context (organization details, competitor names, market data) are sent to Perplexity's API for processing. Uploaded document content is never sent to Perplexity — see File Analysis below.

Perplexity's models have web search enabled by default. This means your queries are used to search the web via Perplexity's own search engine and index as part of their search-augmented generation process.

Expert Search

When you ask the Platform to find experts for a market, the search sentence (role, industry, market, and the business context it is drawn from) is sent to **Exa**, a web-search provider, which returns matching public professional profiles and their LinkedIn URLs. Our older expert page and internal tools use **EnrichLayer** (formerly ProxyCurl) for public-profile lookup. No uploaded document content is included in these searches.

Document Parsing

We use **LlamaParse** (by LlamaIndex) to extract text from uploaded PDFs, Word documents, and other structured formats; plain-text files are decoded on our own servers. LlamaParse caches files for 48 hours and then permanently deletes them. Your data is never used for model training. For details, see [LlamaParse Privacy](#).

File Analysis

For AI analysis of proprietary uploaded documents, we use **Amazon Nova** models through **Amazon Bedrock**, running inside our own AWS account — not Perplexity. Requests go through our AWS endpoint in us-east-2 (Ohio), the region that hosts our database; Bedrock may run the model in other US regions through cross-region inference, never outside the United States. Amazon Bedrock does not store or log prompts and completions, does not use them to train any model, and does not share them with model providers. For details, see [Amazon Bedrock Data Protection](#).

The text of your uploaded files is never sent to Perplexity. Every feature that reads your documents — analysis generation, financial models, and the in-app chat assistant's file questions — routes that text to Amazon Bedrock only. As a second control, our code marks uploaded-document text wherever it is placed in a prompt, and an automated guard rejects any outbound request to Perplexity that carries those markers. This is enforced in code, not only by policy.

One boundary to be clear about: when a user chooses to add insights from a file into their analysis, that analysis text becomes part of the context the Platform works from, including for web research and the chat assistant. What reaches Perplexity in that case is the analysis the user built and can see, never the file itself.

AI Training

We do not use your data to train AI models — ours or anyone else's. We do not operate our own models.

Perplexity publishes a **zero data retention policy** for its API: it does not retain data sent through the API and does not use customer data to train its models. The only data it retains is billing metadata (token counts, timestamps) — never prompt or response content. Perplexity holds SOC 2 Type II certification. For details, see [Perplexity's Privacy & Security documentation](#).

AI Output Disclaimer

AI-generated content on the Platform is for informational purposes only. It may be incomplete, inaccurate, or out of date. Do not treat it as legal, financial, regulatory, or professional advice. You are responsible for verifying any AI-generated output before acting on it.

5. Who We Share Data With

We share data with third-party service providers ("sub-processors") strictly to operate the Platform. Here is the complete list:

PROVIDER	WHAT THEY DO	DATA THEY RECEIVE
AWS (us-east-2, Ohio)	Database hosting (RDS + EKS) and AI processing of uploaded documents (Amazon Bedrock, Amazon Nova)	All Platform data, including uploaded document content
Vercel	Application hosting, file storage, analytics	Application traffic, uploaded files, page views
Perplexity AI	AI analysis and generation with web search	Prompts, organization context (company, product, market, and competitor details), chat messages, and analysis content. Never uploaded document content.
LlamaParse (LlamaIndex)	Document text extraction	Uploaded files (PDF, Word, and other structured formats) and their filenames
Stripe	Payment processing	Names, emails, payment amounts
Hunter	Signup email verification and profile pre-fill	Signup email address; returns name and job-title suggestions
Resend	Transactional email delivery	Names, emails, message content
Exa	Expert search (web search for public professional profiles)	Expert search sentence (role, industry, market, business context); returns public profiles and LinkedIn URLs
EnrichLayer (formerly ProxyCurl)	Public professional-profile lookup (older expert page and internal tools)	Search criteria, LinkedIn URLs
Google Workspace	Slides and Sheets export (Drive)	Generated report and financial-model content. Exported files are shared with your email address. Market-analysis presentations are downloaded through our application; spreadsheet exports and older presentation exports are readable by anyone holding the file link.
Novu	In-app and email notifications	User ID, email address, notification content
Dash0	Application monitoring and error tracking	Logs, traces, performance metrics (logs may include generated document titles and descriptions)
Logo.dev	Company logo retrieval	Company domain names

PROVIDER	WHAT THEY DO	DATA THEY RECEIVE
Inngest	Background job processing	Job payloads and step results, including extracted document text and the uploader's identity while a file is being processed
Slack	Internal team notifications	Support requests, system alerts, call-booking requests, and payment notifications (contact details, organization name, company profile, short excerpts of onboarding answers, subscription tier, and payment amounts)

Link-outs (not sub-processors)

Some buttons take you to a third-party site where you enter information directly. These services receive nothing from our servers.

SERVICE	WHAT IT DOES	DATA YOU PROVIDE
Apollo.io	Meeting scheduling for "Book a call" buttons	Name and email, entered by you on Apollo's booking page. Nothing is sent from our servers.

Public data sources

To build analyses we also query public data services such as the World Bank API (country indicators) and an open exchange-rate feed. These requests carry only country codes or currency codes — never your personal data or content.

We do not sell your personal data. We do not share it with advertisers.

If we add a new sub-processor that handles personal data, we will update this list. Enterprise customers under a DPA will receive 30 days' advance notice of sub-processor changes.

6. Where Your Data Lives

All data is stored in the United States:

- **Primary database:** Self-hosted Gel (PostgreSQL) on AWS in us-east-2 (Ohio)
- **AI processing of uploaded documents:** Amazon Bedrock via our us-east-2 (Ohio) endpoint, inside our own AWS account (US cross-region inference)
- **Uploaded files:** Vercel Blob Storage (US), on a private store where every download is authorized by our application.
- **Application hosting:** Vercel (US)
- **Monitoring:** Dash0 on AWS in us-west-2 (Oregon)

We do not currently offer data residency in other regions.

For users in the EU, UK, and Switzerland, we rely on the **Data Privacy Framework (DPF)** as the legal mechanism for international data transfers. Global Copilot LLC complies with the EU-U.S. Data Privacy Framework, the UK Extension to the EU-U.S. Data Privacy Framework, and the Swiss-U.S. Data Privacy Framework as set forth by the U.S. Department of Commerce.

You can verify our participation on the [Data Privacy Framework List](#).

In compliance with the DPF Principles, we commit to resolve complaints about your privacy and our collection or use of your personal information. Our independent recourse mechanism for DPF inquiries or complaints is [BBB National Programs DPF Services](#).

7. How Long We Keep Your Data

Your data is retained while your account is active. When you close your account, or ask us to, we delete your account and content within 30 days.

Encrypted database backups are retained for 35 days and then rotate out automatically.

We retain billing records (invoices, payment history) as required by tax law — typically 7 years.

8. Your Rights

If you're in the EU/UK (GDPR): You can access, correct, delete, restrict, or port your data. You can object to processing and withdraw consent. You can lodge a complaint with your local supervisory authority.

If you're in California (CCPA/CPRA): You can know what we collect, request deletion, correct inaccuracies, and opt out of sale or sharing. We do not sell personal data. We will not discriminate against you for exercising your rights.

Everyone: Email hello@globalcopilot.com with your request. We respond within 30 days (45 days if we need an extension, and we'll tell you why).

9. Security

We take security seriously, and we'll be straightforward about where we are.

What we do:

- AES-256 encryption at rest for all database storage (AWS KMS)
- AES-256 encryption at rest for uploaded files (Vercel Blob)
- TLS 1.2+ in transit across all connections, HTTP redirected to HTTPS, HTTP Strict Transport Security enabled
- Role-based access controls (Owner, Admin, Member) — see the [Security & IT Overview](#) for the full role matrix
- Database accepts only authenticated, TLS-encrypted connections

- Automated daily backups with 35-day point-in-time recovery
- HttpOnly, Secure, SameSite session cookies
- Rate limiting on authentication endpoints

What we don't have yet:

- SOC 2 Type II certification
- ISO 27001 certification

SOC 2 Type II is on our enterprise roadmap, and we are happy to walk your team through our current controls and timeline. See the [Security & IT Overview](#) for the detail.

10. Cookies

We keep cookies minimal.

Essential (always active): A single authentication session cookie

(`__Secure-better-auth.session_token`) — HttpOnly, Secure, SameSite=Lax, 7-day expiry, renewed during use.

Functional: A short-lived (1 hour) HttpOnly cookie that carries an onboarding invitation code while you sign up through a partner link, and a 30-day preference cookie that remembers the last market-report tab you opened. Neither contains personal data.

Analytics: Vercel Analytics and Vercel Speed Insights process request metadata (IP address, user agent, page URL) to produce page view and performance data, plus named product events such as which module was opened or when an analysis was launched. These product events carry no names, email addresses, or customer-specific identifiers.

We also use localStorage for a few UI preferences (last admin page visited, search form values, display settings). These never leave your browser.

You can control non-essential cookies through your browser settings.

11. Marketing

We may send occasional marketing emails through Resend. Every email includes an unsubscribe link. You can also email hello@globalcopilot.com to opt out.

12. Children

The Platform is for users 18 and older. We do not knowingly collect data from minors.

13. Changes to This Policy

We may update this Policy from time to time. We will post the updated version here with a new effective date. For material changes, we will notify you by email.