SECURITY

Security & IT Overview

How Global Copilot stores, protects, and processes your data. Written for IT and security reviewers.

Global Copilot LLC
Last updated: September 22, 2026

This page complements our Privacy Policy and General Terms & Conditions. Questions: hello@globalcopilot.com.

At a glance

Application hosting: Vercel (United States)

Database: Gel on Amazon RDS, AWS us-east-2 (Ohio), Multi-AZ

Uploaded files: Private Vercel Blob storage (United States)

Document AI: Amazon Bedrock (Amazon Nova) in our own AWS account, Ohio endpoint with US cross-region processing

Web research AI: Perplexity API, zero data retention; never receives uploaded file text

Encryption: AES-256 at rest (database, backups, uploaded files); TLS 1.2+ in transit

Authentication: Passwordless email link or code; no platform-enforced second factor; SAML and SCIM not enabled

Backups: 35-day point-in-time recovery for the database

Assurance: No SOC 2 Type II report, ISO 27001 certification, or independent penetration-test report yet

Sub-processors: 15 providers, listed with what each receives in Privacy Policy section 5

1. Data Categories and Processing Flows

The Platform handles three categories of data:

- **Account information.** An email address is required to create an account; completing signup also asks for the user's name and job title and for basic company details.
- **Business context and public research.** The Platform builds a market analysis for your company from public information. Users can then edit the company, product, and target-customer descriptions, add or remove metrics, edit generated text, add comments, and ask the AI assistant questions whose answers they can add to the analysis. This content is the working context for web research and the assistant.

- **Uploaded documents (optional).** Users can upload their own files to add context, choosing which information from a file to include and where. No upload is required to produce any output.

How uploaded documents are processed

1 **Storage**

Files are written to private, access-controlled US cloud storage (Vercel Blob, encrypted at rest). Downloads are served through our application, which checks that the requesting user belongs to the organization that uploaded the file; documents attached to a specific project additionally require an Owner, Admin, or Project Admin role.

2 **Text extraction**

PDF, Word, and other structured formats are sent to LlamaParse (LlamaIndex), which extracts the text, caches the file for 48 hours, and then deletes it permanently. Plain-text files are decoded on our own servers.

3 **AI analysis**

The extracted text is analyzed by Amazon Nova models through Amazon Bedrock, running inside our own AWS account via our Ohio endpoint. Bedrock may execute in other US regions through cross-region inference, never outside the United States. It is not sent to a public consumer AI service.

4 **Never sent to web search**

The text of your uploaded files is never sent to Perplexity, the web-search model we use for public market research. Routing is the primary control: every feature that reads your documents sends that text to Amazon Bedrock. 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.

5 **User-approved insights**

When a user chooses to add insights from a file into their analysis, that analysis text becomes part of the business context above, which is used for web research and the assistant. What reaches Perplexity in that case is the analysis the user built and can see, never the file itself.

6 **No training**

No customer content is used to train any model, by us or by our providers.

2. Storage, Regions, Encryption, and Retention

All customer data is stored and processed in the United States.

DATA	LOCATION
Application hosting	Vercel, United States
Primary database	AWS us-east-2 (Ohio)
Uploaded files	Vercel Blob, United States
Document AI endpoint	Amazon Bedrock, us-east-2 (Ohio)
Document AI execution	US regions only, through Bedrock cross-region inference
Monitoring and logs	Dash0 on AWS us-west-2 (Oregon)

- **In transit:** TLS 1.2 or higher on every connection (TLS 1.3 in use), HTTP redirected to HTTPS, HTTP Strict Transport Security enabled.
- **At rest:** AES-256 encryption for the database and all of its backup snapshots (AWS KMS) and for uploaded files (Vercel Blob).
- **Database:** Gel on Amazon RDS for PostgreSQL, Multi-AZ with automatic failover. The database accepts only authenticated, TLS-encrypted connections.
- **Secrets:** production credentials live in the hosting platform's secret store, never in source code.
- **Browser protections:** anti-framing (X-Frame-Options: DENY) and MIME-sniffing protection headers; a Content Security Policy is deployed in report-only mode while we tune it.
- **Retention and deletion:** active content is retained while the account is open. When an account is closed, or on request, we delete the account and its content within 30 days. Database backups rotate out after 35 days. Billing records are kept as long as tax and accounting law requires. Details are in Privacy Policy section 7.

3. Authentication and Session Management

Global Copilot is **passwordless**: there is no Platform password to steal, reuse, or brute-force. On the sign-in screen the user enters their email address, and we send an email containing a magic link and a verification code; the user clicks the link or enters the code. Sign-in security therefore rests on the user's mailbox, which is typically protected by your organization's identity provider and MFA policy. The Platform does not enforce a second factor of its own.

- **Link and code validity:** 5 minutes.
- **Sessions:** expire after 7 days and are periodically renewed during use. Cookies are HttpOnly, Secure, and SameSite=Lax.
- **First session:** a brand-new self-serve signup receives its first session immediately, before the email address has been proven by a login. Every subsequent sign-in requires the emailed link or code.

- **Rate limits:** sign-in 10 requests per minute and magic-link requests 5 per minute per client; login-code attempts are limited per window.
- **Federation and provisioning:** single sign-on (SAML) and provisioning (SCIM) are not enabled today. Our authentication system has plugins for both, and we are willing to implement them depending on the agreement we come to.

4. Access Control and User Lifecycle

Access is role-based at two levels. The person who creates an organization is its **Owner**. Every other user holds an organization role (Owner, Admin, or Member) and, where relevant, a project role (Project Admin or Project Member). Organization Owners and Admins hold their permissions across every project in the organization. A user can also be invited to a single project without an organization role; they see that project only.

Organization-role permissions

ACTION	OWNER	ADMIN	MEMBER
View all projects in the organization	Yes	Yes	Yes
Download documents attached to a project	Yes	Yes	Only as a Project Admin
Invite colleagues to the organization (always as Member)	Yes	Yes	Yes
View and manage members on the Team page	Yes	Yes	No
Change member roles	Yes (a remaining Owner is required)	Members and Admins only; cannot promote to Owner	No
Promote another member to Owner	Yes	No	No
Remove members from the organization	Yes (any member except themselves)	No	No
Open the Stripe billing portal	Yes	Yes	No
Archive or restore projects	Yes	Yes	Only as a Project Admin
Manage project members and invite to a project	All projects	All projects	Only as a Project Admin

Project-role permissions

ACTION	PROJECT ADMIN	PROJECT MEMBER
View project data	Yes	Yes
Download documents attached to the project	Yes	No
Archive or restore the project	Yes	No
Manage project members (add, remove, change roles)	Yes	No
Invite users to the project	Yes	No

Invitations

- Any organization member can invite a colleague to the organization; the invitee always starts as a **Member**, and an Owner or Admin can raise the role afterwards. Only Owners, Admins, or Project Admins can invite someone to a specific project.
- Invitation links are valid for 7 days. The email address is not validated on acceptance, so a forwarded link can be accepted by whoever opens it while it is valid.
- Access is email-based: a shared inbox or distribution list can be invited, and anyone who can read that mailbox can use the login link or code.
- Membership does not expire on its own; an Owner can remove a user at any time, and a removed user loses access to the organization and its projects.

Global Copilot staff access

- Staff use the same passwordless login as customers. Administrative functions are limited to accounts on our company email domain and are separate from customer organization roles.
- Staff access customer data only to operate and support the Platform. We do not sell personal data.

5. Third Parties and External Sharing

The complete sub-processor list, with what each provider receives, is maintained in Privacy Policy section 5. The processing boundaries that matter most to reviewers:

- **Amazon Web Services:** database hosting and AI processing of uploaded documents (Amazon Bedrock, Amazon Nova) inside our own account, Ohio.
- **Perplexity AI:** public market research and the assistant, from prompts, business context, and chat, under its published zero-data-retention policy. Never receives uploaded file text.
- **Exa:** expert search. The search sentence (role, industry, market, and the business context it is drawn from) is sent to Exa, which returns public professional profiles. EnrichLayer (formerly ProxyCurl) provides public-profile lookup on our older expert page and internal tools. No uploaded document content is included.

- **Exports to Google Workspace:** exported files are shared with the requesting user's email address. Market-analysis presentations are downloaded through the application; spreadsheet exports and older presentation exports are readable by anyone holding the file link.
- **Payments:** Stripe processes all payments through Stripe Checkout. We never see or store card numbers.
- Customers under a Data Processing Agreement receive 30 days' notice of sub-processor changes. No provider trains models on your data.

6. Secure Development and Vulnerability Management

- Every production change goes through a pull request with automated linting, type checks, unit tests, and end-to-end tests.
- Security regression checks run in that pipeline: an allowlist of every code path that may call Perplexity (so a new path cannot send data out unnoticed), a registry of every model call, and an inventory of every API route's authorization scope.
- Static code security scanning (GitHub CodeQL) runs on every merge and weekly; automated dependency-update proposals (Dependabot) run on a schedule.
- Free text that users type (company, product, segment, and competitor names, notes, file excerpts) is filtered for known prompt-injection patterns before it is placed in a model prompt, and structured output schemas bound what a model can write back. These filters reduce risk; they are not a guarantee against every malicious instruction.

7. Monitoring and Logging

- Operational logs and traces are collected in Dash0 with secrets redacted; scheduled health checks watch the database and background-job system.
- Application activity (page views, instrumented interactions) is recorded and associated with the signed-in user's account.
- Logs may include generated document titles and descriptions; they do not include uploaded file contents. What telemetry providers receive is listed in Privacy Policy section 5.

8. Business Continuity, Backups, and Recovery

- Automated daily database backups, retained for 35 days.
- Continuous point-in-time recovery to any moment within that window; the data-loss exposure for a database incident is on the order of five minutes.
- Documented disaster-recovery and business-continuity procedures, with a four-hour restoration target after a database failure or data-loss event. These are internal targets, not a contractual SLA.
- Uploaded files are stored separately from the database in Vercel Blob and are not part of database backups.

9. Incident Response and Customer Notification

- We maintain documented incident procedures covering containment, credential rotation, session revocation, recovery, and customer communication.
- Customers under a Data Processing Agreement receive notice of incidents affecting their data on the timeline set in that agreement. Personal-data breach notification is described in our Privacy Policy.

10. Compliance and Independent Assurance

- **Available today:** this overview, our Privacy Policy, a Data Processing Agreement on request, and answers to your security questionnaire.
- **Not yet available:** a SOC 2 Type II report, ISO 27001 certification, or an independent penetration-test report.
- **Roadmap:** SOC 2 Type II. We can walk your security team through our current controls and timeline.
- International transfers (EU, UK, and Switzerland) are covered in Privacy Policy section 6.

11. Common IT Questions

Do you enforce MFA?

No. Sign-in relies on access to the user's mailbox, and the Platform does not enforce a second factor of its own. See section 3.

Can exported files be opened outside our organization?

Market-analysis presentations are downloaded through the application. Spreadsheet exports and older presentation exports are readable by anyone holding the file link. See section 5.

Do database backups include uploaded files?

No. Uploaded files are stored separately in Vercel Blob and are not part of database backups. See section 8.

Can we connect Entra ID (SAML 2.0) or provision users with SCIM?

Not today. Our authentication system has plugins for both, and we are willing to implement them depending on the agreement we come to. See section 3.

Is there a standard contractual SLA?

No, unless expressly agreed in a separate written agreement. Recovery targets are in section 8.

Questions we haven't covered? Email hello@globalcopilot.com and we'll answer directly.