SECURITY & PRIVACY

A job search is private. Vector is built so it stays that way.

Résumés, pay, offers and coaching notes are some of the most personal records a person has. Here’s how Vector protects them, who can see them, and the controls every person has over their own.

Encrypted

TLS in transit, HSTS everywhere; encrypted at rest in the database and file storage

Isolated

Firms can't see each other's clients. Access is checked on every request, by role

No passwords

Sign-in by WorkOS. Vector never stores a password

Not for sale

No selling data, no ads, no advertising cookies; no trackers in the app, and website analytics only if you accept

Not for training

Your content isn't used to train AI models

Yours to take

Download everything, or delete the account, from Settings

How the data is protected

- Vector runs on Render in the United States. The database and internal services sit on a private network; the database accepts no connections from the internet.
- Every connection uses TLS (1.2 or 1.3), and browsers are told to use nothing else (HSTS). Data is encrypted at rest in the database and in file storage (Cloudflare R2).
- Uploaded files sit in a private bucket and are only ever served through the app, after the same permission check as the page showing them.
- Internal services require a shared secret on every call, and any page our document renderer loads is checked against the address it resolves to, so it can't be pointed inside the network.
- The database is continuously backed up, with point-in-time recovery over the last 7 days.

Who can see what

- A job seeker sees their own search. Their coach sees the clients assigned to them, and a firm sees its own clients; one firm never sees another's.
- A headhunter sees only what the candidate chose to share with them, and the candidate's own notes on a role only if they chose to share those too.
- A company sponsoring someone's search (outplacement) sees totals across a group, never the search. With Direct Assistance, a helper from that company sees only what the person chose to share, and only while they choose to.
- Permissions are enforced on the server for every request, by the role the person is acting in at that moment, and re-read from the database each time, so removing someone's access takes effect on their next click.
- A note a coach marks “Private to me”, wherever they write it, is visible only to the coach who wrote it: not other coaches, not their firm, and never the client. Each client's contact details (the emails and phone numbers they entered or bought) are theirs alone; no other user sees them.
- Sessions are server-side: the browser holds a random token, we store only its hash, sessions expire after 24 hours, and signing out ends every session.
- Changes to records are written to an audit log.

AI and your data

- Vector's research and drafting use AI models from Anthropic, OpenAI and Google, and research from Perplexity and Exa, through their business APIs. Each task sends only what it needs.
- We don't train AI models on customer content.
- Some deeper searches can be run by outside research agents (ChatGPT, Perplexity). Anyone can switch those off in Settings, and Vector's own agent does the same work instead.
- Anything an AI writes is shown to the person before it goes anywhere. Vector sends, applies to and shares nothing without their decision.

The controls people have

- Download a complete copy of their data, and delete their account and everything in it, from Settings.
- See exactly who can see their search. Nobody can unless the person connects their own coach, a firm invited them, they add a headhunter, or they choose a company helper; a connected coach, a headhunter or a helper can be removed at any time.
- Keep research inside Vector, turn off emails and notifications, and revoke any link they've shared.
- People who aren't users (a contact, or the reader of a shared page) can ask what we hold and have it removed.

How we build it

- Every change runs through the automated test suite, including a dedicated set of tests that check who can and can't reach each kind of record.
- Before launch we ran our own audit of authentication, authorization and tenant isolation, data handling, third parties, dependencies and secrets, and fixed what we found. There were no leaked secrets in the codebase or its history.
- Security headers on every response forbid framing and restrict what pages may load. Production doesn't publish API documentation.
- Secrets live in the hosting platform's encrypted store, never in code.

Questions firms ask

Are you SOC 2 certified?

Not yet. We haven't had an independent security audit. We're glad to complete your security questionnaire or walk your team through this page.

Do you sign a data processing agreement?

Yes. Your firm controls the records it keeps while coaching a client, and we process them on your instructions; we'll sign a DPA covering them. Your client's own account is theirs, under our terms. Ask us at info@vectorjsos.ai.

Is our clients' data used to train AI?

No. The AI providers we use receive what a task needs through their business APIs, and we don't train models on it.

What happens to a client's data when they leave?

Their account is theirs, and stays theirs after your coaching ends. Your firm can download its records about a client, delete the ones it kept for itself, end the coaching, and choose how long its records are kept afterward. The client can download their data and delete their account at any time. Deletion closes the account and signs them out everywhere immediately; their data is then erased from the database and file storage, normally within minutes and always within a day, and from backups within 7 days, along with their sign-in at WorkOS and their customer record at Stripe. Stripe keeps the invoices the law requires, and AI providers keep what they received for a short period under their own terms.

Where is the data stored?

In the United States. We don't offer EU hosting today. The providers we use and where they are are listed below.

What if something goes wrong?

We investigate straight away, contain it, and notify affected customers without undue delay and as the law requires, with what happened and what we're doing about it.

Who receives data

The companies that process personal data for Vector, what each receives, and where. Nothing is sold, and nothing is shared for advertising.

PROVIDER	WHAT FOR	WHAT IT RECEIVES	WHERE
HubSpot	Analytics on our public website (vectorjsos.ai), only if you accept it there; never in the Vector app	The pages you view on the website, where you came from, your IP address (used to estimate the company network you visit from), and a cookie that recognizes your browser	United States
Render	Hosting and the database	Everything Vector stores	United States (Oregon)
Cloudflare	File storage (R2) and network delivery	Files you upload, documents Vector generates, and web traffic	United States / global network
WorkOS	Sign-in	Your name and email address	United States
Stripe	Payments	Your email and payment details (entered with Stripe, never stored by Vector)	United States
Resend	Email delivery, and receiving the newsletters you forward to Vector	Your email address, the emails Vector sends you, and the emails you forward to your Vector address	United States (AWS us-east-1)
Anthropic	AI models (Claude)	The parts of your search a task needs: résumé text, profile, strategy, pay expectations, contacts' names; for an introduction-path search, who you want to reach and the affiliations you list	United States

PROVIDER	WHAT FOR	WHAT IT RECEIVES	WHERE
OpenAI	AI models, job discovery, the ChatGPT research agent, and speech for interview rehearsal	The same, for the tasks it runs; research briefs (name, target roles and companies, locations, pay range) unless you keep research inside Vector; the recording of a rehearsal answer you give out loud, to turn it into text, and the interviewer's questions, to read them aloud	United States
Google	AI model (Gemini) for job discovery	Search queries built from your strategy: titles, locations, exclusions	United States
Perplexity	Web research	Research requests about companies; for deep research, your name, positioning and warm contacts; for an introduction-path search, the name and organization of the person or organization you want to reach	United States
Exa	Web search	Search queries and company names	United States
Coresignal	Organization research you buy, and company measures for How they operate	The company and role filters you choose; for How they operate, only the name and website of a company in a search, never who is searching	EU (Lithuania-based)
Apollo	Finding recruiters at an agency in your search when other sources find too few, and company lookups	Company names and domains, and the job titles to look for; never who is searching	United States

PROVIDER	WHAT FOR	WHAT IT RECEIVES	WHERE
FullEnrich	Organization research and work-email lookup you buy, and looking up where your contacts have worked when you ask it to (on paid plans)	The company and role filters you choose, the names and LinkedIn URLs of the people at that company, and the LinkedIn URLs of the contacts Vector looks up for you	EU (France-based)

The full detail is in our Privacy Policy and Terms of Service. Last reviewed October 2026.