

Privacy Policy

What personal data unused.market collects, why, who receives it, and what you can do about it.

EFFECTIVE

7 October 2026

OPERATOR

Dominik Wójtewicz
Al. Solidarności 68/121
00-240 Warszawa, Poland

CONTACT

contact@unused.market

This policy explains what personal data unused.market collects, why, who receives it, and what you can do about it. It covers the website and dashboard at unused.market and the API gateway at api.unused.market (together, the "Service").

01 Who is responsible for your data

The controller of your personal data is **Dominik Wójtewicz**, a natural person operating unused.market as unregistered activity (*działalność nierejestrowana*) under Article 5 of the Polish Entrepreneurs' Law, Al. Solidarności 68/121, 00-240 Warszawa, Poland ("we", "us").

For anything in this policy, including exercising your rights, write to contact@unused.market. We have not appointed a data protection officer, because the law does not require one for what we do; the address above reaches the person responsible directly.

This policy applies to visitors of the website, to account holders (buyers and sellers), and to people who write to us. It does not apply to the AI providers, payment services and other third parties described in section 5: each of them handles data under its own policy.

Definitions used below. A *buyer* sends requests through the gateway and pays for them with credits. A *seller* connects an API key or subscription from an AI provider and is paid for the requests it fulfils. *Request content* means everything in the body of a request and its response: prompts, messages, system instructions, images, files, tool definitions, tool calls and results, and model output. *Request metadata* means the facts about a request that are not its content.

02 What the Service does, in short

unused.market sells access to AI models through one API key. Sellers connect an API key or subscription from an AI provider and list capacity on it. Buyers add credits, create an unused.market API key, and send requests to our gateway; we route each request to a seller's listing and forward it to the AI provider behind it. Buyers buy from us and sellers are paid by us; buyers and sellers never deal with, or learn about, each other. The same account can buy and sell.

This matters for privacy in one specific way, described in section 5: **the content of a buyer's requests is processed by a third-party AI provider, under a seller's account with that provider.**

03 Data we collect

3.1 Account

- **Email address and password.** The password is stored only as an Argon2id hash; we cannot read it.
- **Notification preferences** (low balance, payout sent, listing sold out, key near its limit).
- **A seller handle:** a random identifier in the form `seller_` plus 12 characters, generated when the account is created. It is not derived from your email.

We do not ask for your name, phone number, postal address or date of birth to open an account. We do not verify your identity and do not collect identity documents.

Your email address also appears as the "actor" in your own audit log, and is used to send the service emails described in section 3.5.

3.2 Sessions and security records

- **Session records:** when you sign in we store the IP address and browser user-agent string for that session, with its creation and expiry time. The session token itself is stored only as a hash.
- **Audit log:** account-level events (account created, sign-in, password or email changed, API key created, changed, revoked or first used, listing created, changed, paused or deleted, routing changed, payout method changed) together with the time, the acting email or key name, and the IP address. You can see and export this log under Analytics → Audit Log.
- **Sign-in rate limiting:** a counter keyed to a hash of your IP address and a hash of the email entered, kept for 15 minutes. It also counts attempts against email addresses that have no account.
- **API rate limiting:** a short-lived counter per API key, kept for about a minute.
- **Application logs:** for each HTTP request, the route template, method, status code and duration. They contain no IP address, no request content, no cookies and no credentials. The web server in front of the application keeps no access log. Logs are kept for troubleshooting for no longer than 30 days.

3.3 Buyer data

- **API keys you create:** the key's name, its last four characters, monthly limit, allowed models, creation time and last-use time. The key itself is shown to you once and stored only as a SHA-256 hash.
- **Routing settings:** strategy, price cap, and your blacklist or whitelist of models.
- **Request metadata:** for every request sent through the gateway we record a request ID, time, model, which API key was used, which listing served it, input, output and cached token counts, the prices applied, cost, latency, status (success, error or rate limited), the HTTP status code the provider returned, the API format used (OpenAI or Anthropic), and whether the token count was reported by the provider or estimated by us. This is what you see under Analytics → Requests. The IP address a request came from is not stored with the request; it is recorded once, in the audit log, the first time a key is used.
- **Holds:** while a request is running, an amount is reserved against your balance and your key's monthly limit. The reservation is deleted when the request settles or after it times out.
- **Credit ledger:** top-ups and usage charges, and daily usage totals per model.

- **Top-ups:** amount, method (card or USDC), status and a payment reference. For card payments this is the payment processor's reference; **we never receive or store your card number**. For USDC payments we store the deposit address generated for your top-up and the on-chain transfers received at it (transaction hash, amount, block number).
- **Payment safeguards:** to stop a top-up or payout from being executed twice, we keep the idempotency key your browser sent with the operation, a hash of the operation and its result.

3.4 Seller data

- **Provider credentials:** the API key or subscription credential you connect. It is encrypted with AES-256-GCM under a per-credential key, which is itself wrapped by a master key held outside the database. It is decrypted only in memory, at the moment a request is forwarded or your available models and balance are checked. It is never returned by the Service, never shown in the dashboard and never written to logs.
- **Source and listing details:** the provider or plan, the models detected on your credential, prices, amounts on offer, tokens sold and status.
- **Model detection:** when you enter a credential, we send it to the provider once to list the models it can call. A credential that is only checked this way, without creating a listing, is not stored.
- **Provider balance estimates:** to show buyers how much capacity a listing really has, we periodically query the provider for the remaining balance on your credential. Raw balances are cached for about a minute, or for up to 15 minutes if the provider is temporarily unreachable, under a fingerprint of the credential rather than the credential itself, and are not shown to anyone.
- **Fulfilment records:** which of your listings served which request, with token counts and the amount owed to you. You do not see who the buyer was.
- **What your provider learns:** requests fulfilled through your listing reach your provider authenticated as you. Your provider therefore sees their content and volume as usage on your account. Your provider may be able to infer from the traffic that it does not all come from you.
- **Earnings and payouts:** your earnings ledger, payout history, and your payout method: either a wallet address (USDC on Base), or the account holder name and IBAN for a payout in fiat currency. For each payout we keep its amount, method, status, a shortened form of the destination, the destination details used, and for USDC payouts the transaction that carried it.

3.5 Emails we send, and messages you send us

We send only service emails, in plain text, with no tracking pixels and no tracked links: a low-balance warning, a notice that a payout was sent, a notice that a listing sold out, and a notice that an API key is close to its monthly limit. Each can be switched off under Account. We keep a record that each notice was sent (its type, time and status) so that it is not sent twice. We send no marketing email and no newsletter.

If you write to us, we keep your message, your email address and our reply for as long as it takes to deal with the matter, and afterwards for up to 3 years in case of a dispute.

3.6 Data we receive from others

- **From Stripe:** confirmation that a card payment succeeded or failed, with its reference and amount. We never receive card numbers, and we do not store the name or billing address you give Stripe.
- **From the Base blockchain:** transfers arriving at the deposit address created for your top-up.
- **From AI providers:** for sellers, the list of models and the remaining balance on the connected credential; for every request, the token usage and status the provider reports.

We do not buy data, do not use data brokers, and do not combine your account with data from social networks or advertising platforms.

3.7 Whether you have to give us data

You need an email address and a password to have an account, and without an account you cannot buy or sell. To be paid as a seller you must give a payout destination. IP address and user agent are sent by your browser or client automatically and are needed to keep sessions secure. Everything else is optional or created by your own use of the Service. You can browse the public marketplace without an account; in that case we store nothing about you.

3.8 Request content, and who is responsible for it

Request content passes through the gateway in memory. We convert it between API formats where needed, count its tokens, forward it to the provider and stream the response back. We do not read it, do not store it, do not log it and do not use it for any purpose of our own.

If request content includes personal data, the buyer who sends it decides why and how that data is processed and is responsible for it under data protection law. We act only as the technical channel that carries it to the provider the buyer's routing selects. Because we keep no copy, we cannot retrieve, correct or delete request content after the fact, and we cannot answer a request for access to it; such requests have to go to the buyer or to the provider.

If your organisation needs a written data processing agreement before sending personal data through the Service, write to us first. Do not send special categories of personal data (health, biometric, genetic, political, religious, trade-union, sex-life or orientation data), data about criminal convictions, payment card data, or data about children.

3.9 Special categories and profiling

We do not ask for, and do not knowingly collect, special categories of personal data. We do not build marketing profiles, do not score users, and do not infer characteristics about you. The only automated processing of your data is described in section 4.

3.10 What we do not collect

- **We do not store the content of requests or responses.** Prompts, messages, images, tool calls and model outputs pass through the gateway in memory so they can be converted between API formats, counted and forwarded. They are not written to our database or our logs. Our application logs record only the route, HTTP method, status code and duration of each request; request bodies, authorization headers and cookies are excluded.
- We run no advertising, analytics or tracking scripts on the Service, and we do not sell or rent personal data.
- We do not use your data to train AI models. We operate no models. Whether an AI provider trains on request content is decided by that provider's own policy (Annex A), not by us.
- We do not use fingerprinting, cross-site tracking, advertising identifiers or session recording, and there is nothing for a "Do Not Track" or Global Privacy Control signal to switch off.
- We do not collect precise location. An IP address reveals an approximate location; we store the address itself and do not derive location from it.

04 Why we use it, and on what legal basis

Where the GDPR or UK GDPR applies, we rely on the following bases.

Purpose	Data	Legal basis
Creating and running your account, signing you in	Account data, sessions	Performance of a contract (Art. 6(1)(b))
Routing and forwarding requests, enforcing key limits	API keys, routing settings, request metadata, seller credentials	Performance of a contract
Billing buyers, crediting sellers, paying out	Ledger, top-ups, payouts, payout method	Performance of a contract
Service emails you have switched on	Email, notification preferences	Performance of a contract; you can switch each off under Account
Security, abuse and fraud prevention, rate limiting, the audit log	IP address, user agent, audit events	Legitimate interests (Art. 6(1)(f)) in keeping the Service and its users' accounts secure
Public marketplace figures (prices, available tokens, 24-hour traffic per model)	Aggregated request metadata, listing details	Legitimate interests in operating a transparent marketplace
Keeping financial and tax records	Ledger, top-ups, payouts	Legal obligation (Art. 6(1)(c))
Handling legal claims and requests from authorities	Whatever is relevant	Legitimate interests; legal obligation

Legitimate interests, weighed. Where we rely on legitimate interests, we have considered whether they are overridden by yours. Keeping IP addresses with sessions and audit events lets you and us detect account takeover; it is limited to sign-in and account changes, is visible to you in your own audit log, and is deleted on the schedule in section 7. Public marketplace figures are aggregated per model and show no individual buyer. You may object to either at any time (section 10).

No consent-based processing. We currently do nothing that relies on your consent. If that changes, we will ask first, and you will be able to withdraw consent as easily as you gave it.

No other purposes. We do not use your data for advertising, do not sell it, and do not share it for anyone else's marketing. If we ever need to use data for a purpose that is not compatible with those in the table, we will tell you beforehand.

We make no decisions about you that are based solely on automated processing and produce legal or similarly significant effects. Automated routing chooses which listing serves a request, based on price, speed and your routing settings; a listing may be paused automatically if its provider credential is rejected, and a request is refused automatically when credits or a key's monthly limit run out or a rate limit is reached. You can ask us to review any of these by writing to contact@unused.market.

05 Who receives your data

5.1 AI providers and sellers

When a buyer sends a request, the gateway forwards its content to the AI provider behind the listing that serves it (for example Anthropic, OpenAI, Google, DeepSeek or another provider on the marketplace), authenticated with **the seller's credential**. This means:

- The AI provider receives the full content of the request and produces the response. It processes that content as an independent controller, under its own terms and privacy policy (each provider's policy is linked in Annex A), and may be located outside your country. Which provider serves a request depends on the model you call and on your routing settings.
- From the provider's side, the request belongs to the seller's account. Depending on the provider, the seller may be able to see usage figures, and in some cases request logs, in their own provider dashboard. **We do not control what a provider makes available to its account holder.**
- We do not send the provider your email address, your IP address or your unused.market API key. If your own client sets an end-user identifier in the request (`user` in the OpenAI format, `metadata.user_id` in the Anthropic format), it is forwarded as part of the request.
- **Sellers** never receive buyer identities or request content from us. **Buyers** see only a seller's handle, never their email or credential.
- A request is sent to one provider at a time. If that provider fails before any part of the response has been returned, the gateway may retry the same request with another listing, which can mean another provider. Your routing blacklist or whitelist limits which models, and so which providers, can be used.
- Some providers are themselves gateways that pass requests on to further providers under their own policies (see the note under Annex A).
- We have no contract with AI providers about your data and cannot give you assurances on their behalf about retention, human review, training or location. Read the policy of each provider you allow in your routing.

Do not send through the Service any content you are not permitted to share with a third-party AI provider, and avoid sending special-category or other highly sensitive personal data. If you send personal data about other people, you are responsible for having a lawful basis to do so.

5.2 Other users and the public

Listings are public: model, provider, price, available tokens and the seller handle. Aggregate traffic figures per model are public. Nothing else in your account is visible to other users.

5.3 Service providers

Recipient	What it receives	Why
Stripe (https://stripe.com/privacy)	Top-up amount, an internal reference, and whatever you enter on Stripe's checkout page	Card payments. Stripe handles your card details as its own controller.
The Base blockchain and our blockchain node (RPC) provider	Deposit and payout wallet addresses, amounts	USDC top-ups and payouts
Peer.xyz	The payout details you gave for a fiat payout, and the amount	Paying out in fiat currency. Peer.xyz is a peer-to-peer service that operates under its own terms and privacy policy.
Our email delivery provider	Your email address and the notification text	Sending service emails
MatHost.eu (https://mathost.eu/documents/privacy-policy/)	All data stored by the Service	Hosting the application, database and cache on servers in Warsaw, Poland

Service providers that process data for us may use it only to provide their service to us. Stripe and Peer.xyz also act for their own purposes, such as fraud prevention and legal compliance, under their own policies. We do not give any of them request content, and none of them receives your password or seller credentials in readable form.

Blockchain transactions are public and permanent. If you top up or get paid in USDC, the addresses, amounts and times are recorded on Base, visible to anyone, and cannot be altered or erased by us. Blockchain analysis can link addresses to one another and sometimes to a person. If that matters to you, use a wallet address you are comfortable having associated with these payments.

5.4 Authorities and successors

We disclose data to courts, law enforcement, tax and other authorities only where a law or a binding order requires it, and only the data that is asked for. Because we do not store request content, we cannot hand it over. Where the law allows, we tell the account holder about a request before answering it.

We may also use and disclose data where it is necessary to establish, exercise or defend legal claims, to investigate fraud or abuse of the Service, or to respond to a provider or credential owner who reports that a credential is being used without authority.

If the Service is transferred to a company set up to operate it, or is merged, acquired or sold, your data will be transferred to the new operator, who will remain bound by this policy until it is changed in line with section 12. We will tell you by email before that happens.

5.5 What we never do

We do not sell personal data, do not rent it, do not exchange it for anything of value, and do not share it for cross-context behavioural advertising. We have not done so in the past 12 months.

06 International transfers

The Service is hosted on MatHost.eu servers in Warsaw, Poland, so the data we store stays in the European Union. Some recipients in section 5, including AI providers and Stripe, process data in the United States and other countries outside the European Economic Area. Where we transfer personal data out of the EEA or the UK, we rely on an adequacy decision or on the European Commission's Standard Contractual Clauses. For requests forwarded to an AI provider, the destination country is determined by the provider you route to. Providers in Annex A are established in, among other places, the United States, China, Singapore, South Korea, Israel and the European Union, and some of those countries have no adequacy decision from the European Commission. Sending request content to such a provider is a transfer you choose by calling its models; if you must keep data in a particular region, restrict your routing to providers that meet that requirement, or do not send that data.

You can ask us for a copy of the safeguards we rely on for our own transfers by writing to contact@unused.market.

07 How long we keep it

Data	Kept for
Account data	While your account exists
Sessions	Until you sign out or the session expires (30 days after sign-in). Changing your password ends every other session.
Sign-in rate-limit counters	15 minutes

Data	Kept for
Request metadata	12 months, then deleted or reduced to daily totals
Audit log	12 months
Seller credentials	Until you delete the listings that use them or close your account
Payout method	Until you replace it or close your account
Ledger, top-ups and payouts	5 years after the end of the calendar year of the transaction, as required by accounting and tax law
Revoked API keys	The hash, name and last four characters are kept with the request history they belong to
Deleted listings	Hidden at once; the record is kept with the request and payment history it belongs to
Balance holds for running requests	Until the request settles, at most a few minutes
Cached provider balances	About a minute, at most 15 minutes
Payment safeguards (idempotency records) and the log of notices sent	12 months
Application logs	Up to 30 days
Correspondence with us	Up to 3 years after the matter is closed
Backups, if any are kept	Deleted data can remain in a backup for up to 30 days until the backup is overwritten

When you close your account we delete or anonymise your personal data within 30 days, except for the financial records above and anything we must keep to resolve a dispute or comply with the law. Records that stay are kept under an internal identifier, without your email address, wherever the law allows.

Three things cannot be deleted by us at all: transactions already recorded on the Base blockchain, data already received by an AI provider or payment service, and anything you made public or sent to someone else yourself.

Where a period above ends, we delete the data or make it anonymous. Where the law sets a longer period, or where data is needed for a claim that is already under way, we keep it until that reason ends.

08 Cookies and browser storage

We use only what the Service needs to work. There is no cookie banner because there are no optional cookies.

Name	Type	Purpose	Lifetime
<code>unused_session</code>	Cookie (HttpOnly, Secure)	Keeps you signed in	30 days
<code>unused_csrf</code>	Cookie	Protects forms and actions against cross-site request forgery	Until the browser is closed or you sign out

Name	Type	Purpose	Lifetime
<code>unused:intro-seen</code>	Session storage	Remembers that the intro animation has played, so it is not shown again	Until the tab is closed

These are strictly necessary for a service you asked for, so they do not need consent. You can block or delete them in your browser, but you will not be able to sign in without them. We set no cookies for analytics, advertising or preferences, and no third party sets cookies on our pages.

Fonts are served from our own domain; loading a page does not contact Google or any other third party. If you pay by card you are sent to Stripe's checkout page, which sets its own cookies under Stripe's policy. API clients calling the gateway need no cookies at all; they authenticate with an API key.

09 Security

Passwords are hashed with Argon2id. API keys and session tokens are stored only as hashes. Seller credentials are encrypted as described in section 3.4. All traffic uses TLS. State-changing requests require a CSRF token, sign-in attempts are rate-limited, and logs exclude request content and credentials.

In more detail:

- **In transit:** connections to the dashboard and the gateway are encrypted. The session cookie is marked `HttpOnly` and `Secure`, so page scripts cannot read it and it is never sent over an unencrypted connection.
- **At rest:** seller credentials are encrypted with a separate key for each credential, and those keys are wrapped by a master key that is kept outside the database and can be rotated. Passwords, API keys and session tokens are stored as one-way hashes and cannot be recovered, by us or by anyone who obtains the database.
- **Access:** account data is reachable only by the account that owns it. Operational access to the servers is limited to the operator; the hosting provider has physical custody of the machines.
- **Abuse controls:** sign-in is limited per IP address and per email, the gateway is limited per key, and unknown-email sign-in attempts take the same time as real ones so that the Service does not reveal which emails have accounts.
- **Your audit log:** every sensitive change to your account is written to a log you can read and export, with its time and IP address.

If something goes wrong. If we become aware of a personal data breach that is likely to put your rights at risk, we will report it to the supervisory authority within 72 hours where the law requires, and tell affected users without undue delay what happened, what data was involved and what to do. If a seller credential may have been exposed, we will tell the seller so that it can be revoked at the provider.

No system is perfectly secure. Keep your password and API keys private, set monthly limits on your keys, and revoke any key you believe is exposed. We cannot guarantee that unauthorised access will never happen, and nothing in this policy is a warranty of security.

To report a vulnerability, write to contact@unused.market.

10 Your rights

Depending on where you live, you have the right to:

- **access** your data and receive a copy. Much of it is available directly: requests and the audit log can be exported as CSV from the dashboard;
- **correct** it. You can change your email, password, payout method and settings under Account and Payouts;
- **delete** it, by closing your account, subject to the records we must keep (section 7);
- **restrict** or **object to** processing based on legitimate interests;
- **data portability** for data you provided, in a machine-readable format;
- **withdraw consent** where we rely on it;
- **complain** to a supervisory authority. In Poland this is the President of the Personal Data Protection Office (UODO), ul. Stawki 2, 00-193 Warszawa; elsewhere in the EEA, your local authority.

How to use these rights. Write to contact@unused.market from the email address on your account, so that we can be sure the request comes from you; if you write from another address we will ask you to confirm from the account's address. We answer within one month. If a request is complex or we receive many, the law lets us extend that by up to two months, and we will tell you if we do. Exercising your rights is free; we may refuse or charge a reasonable fee only for requests that are manifestly unfounded or excessive.

What we can and cannot do.

- *Access and portability:* we will send you the data held under your account in a common machine-readable format (CSV or JSON). It will not include request content, because we have none, or readable passwords, keys or credentials, because we cannot read them.
- *Deletion:* we delete what the law does not require us to keep (section 7). We cannot delete blockchain records or data held by AI providers and payment services; for those, contact the party that holds the data.
- *Objection:* you can object to processing based on legitimate interests. We will stop unless we have compelling grounds that override your interests or need the data for legal claims. Some security processing cannot be switched off for one account while it remains open; in that case the alternative is closing the account.
- *Restriction:* while a dispute about your data is being resolved, you can ask us to keep it without using it.
- *Correction:* most data can be corrected by you in the dashboard. Ledger entries and the audit log are records of what happened and are not edited; if one is wrong, we add a correcting entry.

Someone else's data. If you believe your personal data was sent through the Service by one of our users, we cannot find or delete it, because request content is not stored. Tell us what you know and we will pass your request to the account holder if we can identify one.

Residents of California and other US states. We do not sell personal data and do not share it for cross-context behavioural advertising. You may ask to know, correct and delete your personal data as described above, and we will not treat you differently for doing so.

11 Children

The Service is for businesses and adults. It is not directed at anyone under 18, and we do not knowingly collect their data. If you believe a minor has opened an account, tell us and we will close it and delete its data, subject to section 7.

12 Changes to this policy

We will post any change here and update the effective date. If a change materially affects how we use your data, we will email account holders at least 14 days before it takes effect. Changes that only add detail, correct an error, name a new service provider of a kind already listed, or update a link in Annex A take effect when posted. Earlier versions are available on request.

If any part of this policy conflicts with a law that applies to you and cannot be set aside, that law prevails and the rest of the policy stays in force. This policy describes how we handle data; it is not a contract and does not create rights beyond those the law gives you. The Terms of Service govern your use of the Service.

13 Contact

Dominik Wójtewicz

Al. Solidarności 68/121, 00-240 Warszawa, Poland

contact@unused.market

Annex A. Privacy policies of AI providers

These are the providers a listing on unused.market can be connected to. A request is forwarded to exactly one of them: the provider behind the listing that serves it. Each processes request content under its own policy, linked below. Links were last checked on 7 October 2026.

Provider	Privacy policy
AI21 Labs	https://www.ai21.com/terms-policies/privacy-policy/
AkashML	https://akashml.com/privacy
Alibaba Cloud Model Studio	https://www.alibabacloud.com/help/en/legal/latest/alibaba-cloud-international-website-privacy-policy
Amazon Bedrock	https://aws.amazon.com/privacy/
Anthropic	https://www.anthropic.com/legal/privacy
Atlas Cloud	https://www.atlascloud.ai/privacy
Azure AI Foundry	https://www.microsoft.com/en-us/privacy/privacystatement
Baidu Qianfan	https://intl.cloud.baidu.com/en/doc/Agreements/s/Plr0fi68q-intl-en
Bankr	https://bankr.bot/privacy-policy
Baseten	https://www.baseten.co/privacy-policy/
Capminal Gateway	No published privacy policy found when last checked
Cerebras	https://www.cerebras.ai/privacy-policy
CheaperInference	https://www.cheaperinference.com/legal/privacy
Chutes	https://chutes.ai/privacy

Provider	Privacy policy
Cloudflare Workers AI	https://www.cloudflare.com/privacypolicy/
Cohere	https://cohere.com/privacy
Concentrate.ai	https://concentrate.ai/legal/privacy-policy
Crusoe	https://legal.crusoe.ai/#cloud-privacy-policy
DeepInfra	https://deepinfra.com/privacy
DeepSeek	https://cdn.deepseek.com/policies/en-US/deepseek-privacy-policy.html
Engy.ai	https://engy.ai/privacy
Featherless	https://featherless.ai/privacy
Fireworks AI	https://fireworks.ai/privacy-policy
Friendli	https://friendli.ai/privacypolicy
GitHub Models	https://docs.github.com/en/site-policy/privacy-policies/github-general-privacy-statement
Google Cloud Vertex AI	https://cloud.google.com/terms/cloud-privacy-notice
Google Gemini API (AI Studio)	https://policies.google.com/privacy and https://ai.google.dev/gemini-api/terms
Groq	https://groq.com/privacy-policy
Huawei Cloud	https://www.huaweicloud.com/intl/en-us/declaration/sa_prp.html
Hugging Face	https://huggingface.co/privacy
Hyperbolic	https://www.hyperbolic.ai/privacy
Inception	https://www.inceptionlabs.ai/docs/privacy-policy
Inference.net	https://inference.net/privacy-policy/
Jatevo, Jatevo API Gateway	https://jatevo.ai/privacy
Kluster	https://www.kluster.ai/privacy-policy (unreachable when last checked)
Lambda	https://lambda.ai/legal/privacy-policy
Meta Llama API	https://www.facebook.com/privacy/policy/
MiniMax	https://platform.minimax.io/protocol/privacy-policy
Mistral AI	https://legal.mistral.ai/terms/privacy-policy
Moonshot AI	https://platform.kimi.ai/docs/agreement/userprivacy
Mordiem	https://mordiem.com/privacy

Provider	Privacy policy
Morpheus	https://app.mor.org/privacy
Nebius	https://docs.nebius.com/legal/privacy
Novita AI	https://novita.ai/legal/privacy-policy
NVIDIA NIM	https://www.nvidia.com/en-us/about-nvidia/privacy-policy/
Ollama Cloud	https://ollama.com/privacy
OpenAI	https://openai.com/policies/privacy-policy/
OpenRouter	https://openrouter.ai/privacy
Parasail	https://www.parasail.io/legal/privacy-policy
Perplexity	https://www.perplexity.ai/hub/legal/privacy-notice
Replicate	https://replicate.com/privacy
SambaNova	https://sambanova.ai/privacy-policy
SiliconFlow	https://docs.siliconflow.com/en/legals/privacy-policy
StepFun	https://platform.stepfun.ai/docs/en/agreement/userprivacy
Targon	https://targon.com/privacy
Tencent Cloud	https://www.tencentcloud.com/document/product/301/17345
Together AI	https://www.together.ai/privacy
TypeSafe	https://typesafe.ai/legal/privacy-policy
Uncensored AI	https://uncensored.ai/privacy-policy
Upstage	https://www.upstage.ai/privacy-policy
Venice AI	https://venice.ai/legal/privacy-policy
Vercel AI Gateway	https://vercel.com/legal/privacy-notice
Volcengine	https://docs.volcengine.com/docs/Legalagreements/VolcanoEnginesPrivacyPolicy (Chinese)
xAI	https://x.ai/legal/privacy-policy
Zai, Zai Coding	https://docs.z.ai/legal-agreement/privacy-policy

Several of these are gateways or marketplaces themselves (for example OpenRouter, Vercel AI Gateway, Hugging Face, Chutes, Morpheus, Targon, Concentrate.ai, Jatevo, CheaperInference, Mordiem, Bankr). A request sent to one of them is passed on again to the model provider it selects, under that gateway's own policy.