NOT ORGANIC

Privacy Policy

What the shared account processes, where it goes, and the choices you should have.

Prepared 5 September 2026 · Review copy

Prepared for review, not yet effective. This is substantive proposed wording. The operator’s legal identity, contact details, applicable jurisdiction, and final policy decisions must be confirmed before adoption.

At a glance

- Your requested features determine what information is processed.
- Hosted AI providers receive the content needed to answer a request.
- Stored gateway responses currently expire after 30 days; other records have different retention needs.
- You can ask about access, correction, consent, and deletion.

This summary is a guide to the full text below.

Who and what this policy covers

This proposed policy covers the Not Organic website and shared account, authorization, billing, and hosted-service layer. The adopted policy will identify the responsible legal entity, public business address, and privacy contact.

Connected products can collect additional information for their own features. Their notices should explain that processing. This page does not claim that every document, message, relationship record, or learning record lives in the shared account layer.

Information processed by the shared service

A payment provider collects the payment details needed for checkout. Not Organic receives transaction and entitlement information rather than handling your full card number through its own checkout form. The current billing integration also sends an account identifier and product information to associate a purchase with the correct account.

Why information is used

The purposes are to deliver requested features, authenticate accounts, apply permissions, maintain balances, process purchases, support users, prevent abuse, investigate incidents, and meet legal obligations. Processing should be limited to information reasonably needed for those purposes.

Where consent is required, it must be meaningful and relate to the processing being proposed. Optional features, especially those involving audio, video, or sensitive information, should explain what will be sent and obtain any required permission before transmission. You can decline an optional feature, though that feature may then be unavailable.

For people protected by laws requiring a specified legal basis, the adopted policy must identify the applicable basis for each purpose, such as performing a contract, consent, legal obligation, or a permitted legitimate interest. It must not treat acceptance of terms as consent to unrelated processing.

AI providers and sensitive content

A hosted AI request passes through the gateway and the selected routing or model provider. Those services need the request content to perform the requested task. The gateway substitutes a pseudonymous usage identifier for direct account identity on its inference path, but the content of a prompt can still identify you or someone else.

Do not assume an AI request is end-to-end encrypted from the service operator or invisible to the selected provider. Encryption at rest and protected transport serve different purposes. Provider retention and training terms depend on the configured provider and agreement; this policy does not promise universal zero retention or “never used for training.”

Only submit sensitive information when necessary for the feature and when you have authority to share it. Product notices should explain any use of information to personalize or evaluate AI behavior, and any separate research or model-improvement purpose, before that use occurs.

Service providers and other disclosures

The architecture uses Railway for hosting, Convex for application records, OpenBao for key operations, and configured routing and model providers for inference. Depending on the enabled feature, it can also use Paddle for purchases, PostHog for operational analytics, object-storage services for artifacts, and execution or video providers such as Sprites or Tavus.

This is an explanation of the service categories and supported integrations, not a statement that every provider receives every user’s data. Provider access depends on the feature, configuration, and information needed to supply it. The adopted policy should link a maintained provider list with processing locations and purposes.

We may also disclose information when you direct us to, where required by law, or where lawfully necessary to protect people, investigate abuse, or defend legal rights. A business transfer involving personal information must preserve applicable privacy obligations and include any required notice.

Data may be processed outside your province or country and be subject to the laws there. The adopted policy must describe relevant transfer safeguards rather than promise local-only storage.

Cookies, preferences, and analytics

The account portal stores a notorganic_appearance preference cookie for up to 30 days across Not Organic subdomains. It remembers only the selected product appearance and light or dark theme, so sign-in can match the product you came from. It does not contain an account identifier or authorize access. You can change the appearance in the portal or clear the cookie in your browser.

The public site’s theme control follows your device preference and keeps a manual change only for the current page. The current theme script does not store a persistent preference or set an analytics cookie.

The account portal uses an essential, opaque session cookie, configured for up to 14 days, and short-lived authorization state. In production, the session cookie is protected with Secure, HttpOnly, and SameSite settings. Blocking essential cookies may prevent sign-in.

When gateway analytics is configured, its event payloads include a pseudonymous account identifier and operational metadata such as usage quantities, costs, and outcomes. The telemetry implementation does not place prompt bodies or direct decentralized identifiers in those analytics events. Pseudonymous information can still be personal information.

Additional optional tracking or marketing must be disclosed and use any consent controls required by law. A browser theme setting is not an analytics-consent setting.

Retention and deletion

Retention depends on the purpose and record type. The following describes the implemented defaults; it is not a promise that every backup or external provider deletes on the same schedule.

Where supported and authorized, deleting an individual stored response removes that gateway record. Account closure is broader and needs its own process. We should explain any information retained, why it is retained, and any limits on deleting independent recipients' copies.

Your choices and requests

You can ask for access to your personal information, correction of inaccurate records, an explanation of processing, withdrawal of consent where applicable, and deletion or portability where the law provides it. Depending on where you live, you may also have rights to object, restrict processing, or appeal a decision.

Use the support channel in your product to request help while the dedicated privacy contact is being finalized. Describe the request and relevant account without sending passwords, private keys, or full payment-card details. Identity checks should be proportionate to the request.

We should respond within the time required by applicable law, explain any lawful refusal or extension, and avoid requiring unnecessary information. A request may affect a feature that depends on the information, but asking to exercise a right should not itself lead to unfair treatment.

You can raise a privacy concern with the authority responsible for your location. In Canada, the Office of the Privacy Commissioner of Canada explains complaint options.

Security and identity portability

The shared layer uses scoped authorization, device-bound proofs, access controls, and encryption for stored response content. These measures reduce risk; they do not eliminate it. We should assess incidents and notify affected people or authorities when required by law.

AT Protocol identity and public records have different visibility from private account and billing records. Public identifiers or records may be discoverable or replicated by independent services. Moving an identity does not automatically move every connected product's private data or erase previously shared public records.

Children and product-specific information

Products serve different purposes and may have different minimum ages. A learning feature and a dating feature cannot be treated as having the same eligibility or privacy needs. The adopted policy and each product's onboarding must state the applicable ages, parental-permission process where relevant, and treatment of children's information.

If you believe a child has supplied information contrary to a product's rules or without required permission, contact the product's support channel so it can be investigated.

Changes and accountability

The adopted policy will carry an effective date and identify a person or team accountable for privacy. Material changes should be communicated appropriately, and a new purpose requiring consent must not be introduced solely by changing this page.

This review copy is based on the service implementation and published privacy principles. The final operator identity, contact, age policy, provider arrangements, retention schedule, and regional rights details still need confirmation.

Online version: https://notorganic.info/privacy