# Privacy Policy

How Wakemark processes data when you route generative-media jobs through the control plane — what we collect, why and on what legal basis, how long we keep it, where it lives, and the rights you have. Written from what the system actually does.

Last updated: 2026-07-13

## 1. Who we are and what this covers

Wakemark is a control plane for generative media: one API and dashboard above the generation providers, with capability routing and cross-provider failover, byte-exact custody of outputs, C2PA provenance preservation and detection, an append-only audit ledger, and a public verifier. This policy covers the Wakemark website, dashboard, and API.

**THE OPERATOR**

The service is operated by O ÷W¦† â ±1r, an individual (a natural person) based in Türkiye. There is currently no registered company behind the service — “Wakemark” is the product and service name, not a company name. For privacy, data-protection, or legal inquiries, contact us at hello@wakemark.ai.

## 2. Our two roles: controller and processor

Wakemark acts in two distinct capacities, and your rights run differently under each. Getting this split right matters more than anything else in this document.

- As controller — for account and workspace data (your email and authentication identifiers, organisation/workspace records, API-key metadata, support correspondence), for website request metadata and aggregate analytics, for transactional email, and for the PII-free erasure receipts described in section 7. Here Wakemark decides the purposes and means, and this policy is the governing document.
- As processor — for customer content: the prompts and parameters you submit, the media that comes back, the job/event ledger entries that record them, C2PA provenance sidecars, and the lineage graph — plus the encrypted custody of the BYOK provider and storage credentials you bring. Here you (or your organisation) are the controller and Wakemark processes on your documented instructions under the Data Processing Agreement.

**THE BYOK REALITY — DATA WAKEMARK NEVER SEES**

Under bring-your-own-key, your relationship with each generation provider is directly yours: your account, your key, their terms, their bill. Wakemark never sees your provider account or its billing data. If you connect your own storage bucket, your generated bytes land there — in a bucket Wakemark can write to only under the credentials you granted, reads only to serve you, and never sweeps or lists beyond the objects it wrote. Your provider keys are stored envelope-encrypted and exist in plaintext only in memory at the moment a job is submitted; they are never written to logs or error messages.

## 3. What we process

**Account & workspace data (controller)**

Email address and authentication identifiers (via our authentication provider), organisation/workspace name and membership, API-key metadata (a hash and prefix — never the raw key, which is shown to you once and stored only as a SHA-256 digest), and project names and settings.

**Customer content (processor — on your instructions)**

Prompts, parameters, and reference inputs you submit; the generated media returned by the providers you route to; the append-only job/event ledger recording each job's lifecycle (including raw provider responses, which can embed your input or output); provenance sidecars (digest, model, provider, C2PA detection result, lineage); and the cross-provider lineage graph. All of it is subject to the retention policy you set (section 6).

**BYOK credentials (processor — encrypted custody)**

The provider API keys and storage-bucket credentials you connect, sealed with per-secret envelope encryption under a managed key-management service (KMS), cryptographically bound to your tenant so ciphertext copied elsewhere is useless. See Security & Trust.

**Operational logs & telemetry (controller)**

Structured request logs carry method, path, status, and tenant — never prompts, media, or secrets. Operational metadata in the ledger (timestamps, provider, model, status, estimated cost) supports the service and may be aggregated into content-free, platform-wide statistics that never expose any tenant's individual traffic.

**Website data (controller)**

HTTP request metadata from hosting, and cookieless aggregate page-view analytics (section 9). No advertising identifiers.

## 4. What we never do

Wakemark will never train on, mine, or repurpose your content — not for models, not for datasets, not in “anonymised” form. We are not a model builder, and we do not sell personal data. What Wakemark measures instead is content-blind operational metadata — latency, cost, and reliability per provider and model — which touches none of your bytes. This is a binding commitment, made identically in the Terms of Service and the DPA.

- No advertising cookies, no cross-site trackers, no profiling.
- No sale or rental of personal data, to anyone, for anything.
- No training on, mining of, or repurposing of customer content — including “anonymised” derivatives.

- No payment or billing data: nothing is charged today. Wakemark has no billing system in the current release, holds no payment-card data, and issues no invoices. The dashboard's storage-cost figure is an estimate, labelled as such, and is never a charge. If usage-based billing launches, this policy and the subprocessor list will be updated first.
- No plaintext provider keys at rest, in logs, or in error messages — ever.
- No silent deletion from, or scanning of, your own storage bucket: Wakemark touches only objects it wrote and recorded, and only to execute the retention policy you set.

## 5. Purposes and legal bases

Where the GDPR applies and Wakemark acts as controller, we rely on the following legal bases under Article 6(1):

- Providing the service you signed up for — accounts, authentication, projects, dashboards, exports (Article 6(1)(b), performance of a contract).
- Transactional notices — e.g. the account-deletion confirmation email (Article 6(1)(b); Article 6(1)(f) where not strictly contractual).
- Security and abuse prevention — tenant isolation, credential protection, rate-limiting the public verifier, investigating misuse (Article 6(1)(f), legitimate interests).
- Service operation and improvement using content-free data — aggregate reliability/latency statistics and cookieless page-view analytics (Article 6(1)(f), legitimate interests).
- Legal compliance — responding to lawful requests and keeping records the law requires of us (Article 6(1)(c)).

For customer content, Wakemark processes on your documented instructions as processor; you are responsible for having your own lawful basis for the content you submit and the people it concerns.

## 6. How long we keep it

Content retention is customer-controlled, per project, with two independent knobs — one for inputs (your prompts and reference material), one for outputs (the generated media) — each settable to none, immediate, a fixed duration, or forever, with a per-job override. New projects default to keeping inputs for 30 days and outputs until you delete them or change policy. When content passes its deadline, a retention worker deletes the bytes and redacts the content-bearing fields in the ledger, leaving an auditable tombstone (digest, size, provider, model, lineage — no content).

- “None” for inputs means no content survives the job's terminal state — not “never touches our systems”: the engine holds the input in the ledger while the job runs (failover re-submits from it) and redacts it the moment the job completes.
- “None” for outputs waives custody entirely: Wakemark never copies the output, and you must fetch it from the provider's own URL before the provider expires it. This opts out of the core custody promise, and the product says so at the point of choice.
- “Immediate” for outputs stores the output for a short retrieval window (48 hours by default) so you can download what you paid your provider for, then purges it.
- A usage-record floor survives content redaction: job id, project, capability, provider/model, timestamp, and estimated cost — no prompt, no media. “Never store” means “never keep my content”, not “erase that I used the product”. This floor is itself deleted if you delete your account.

- If your project stores to your own bucket, the same policy is executed there — deleting only objects Wakemark wrote and recorded, never a sweep of your bucket — and if Wakemark's credentials cannot delete, it fails loudly and records no purge.

## 7. Deleting your account (erasure)

Account deletion is built as a first-class, proven flow — not a support ticket. When you request deletion:

- Immediately (T+0): your stored provider credentials are deleted outright and all API keys are revoked — the assets that could spend your money never sit in a deletion queue. In-flight jobs are cancelled best-effort while the credentials still exist.
- A 30-day grace period (configurable) follows, during which the account is frozen: no generation, no mutations — only cancelling the deletion or downloading a full account export. This window exists to recover from mistaken or coerced deletion.
- At the deadline, a hard purge erases every content store we control: platform-stored media and sidecars, the jobs projection, the event ledger, assets, lineage, API keys, storage-connection credentials, projects, and the tenant record, plus deletion of the associated authentication organisation where configured.
- Your own bucket is never touched. Offboarding deletes our access and our records — your bytes stay yours, untouched.
- What survives: a PII-free erasure receipt — a one-way-hashed tenant identifier, a timestamp, and counts of what was deleted — plus anonymized usage-metering records (job counts, storage byte sizes, timestamps, and the random job/asset identifiers they counted, with account and project references replaced by the same one-way hash). No content, names, or account details: together they prove the purge ran and preserve aggregate usage without keeping any of what was deleted.
- Backups: our database and storage vendors hold their own backups (e.g. point-in-time recovery); a purge does not reach into those, and they age out on the vendor's retention window. We disclose this rather than claim otherwise.

**ERASURE IS TESTED, NOT ASSERTED**

The deletion flow is verified by automated tests that seed a tenant with content and then search every database table and storage object for it after the purge — on both supported database backends. We describe what the system does; we do not decorate it with the word “guaranteed”.

## 8. Where your data lives, and international transfers

- Object storage — platform-stored media and sidecars, in the European Union.
- Database — account, project, and ledger records, in the European Union.
- Key management — envelope encryption of the credentials you bring, in the European Union (Frankfurt region).
- Compute — job execution, in the European Union.
- US-based services — authentication, transactional email, and web hosting + analytics process in the United States. This transfer is covered by the Standard Contractual Clauses under the DPA.

- Generation providers and your own bucket — under BYOK your prompts run on the providers you connect, under your account and their terms, in their regions; bytes you direct to your own bucket live wherever you put that bucket.

## 9. Cookies & analytics

The browser surface is deliberately lean — no advertising cookies, no cross-site trackers.

- Strictly-necessary cookies — signing in sets an encrypted session cookie so you stay authenticated between requests. Without it the app cannot keep you logged in.
- Privacy-friendly analytics — the marketing site uses a cookieless analytics service that measures aggregate page views without profiling you or tracking you across sites. No analytics cookies, no advertising identifiers.

## 10. Your rights

Where the GDPR applies, you have the rights of access, rectification, erasure, restriction, portability, and objection (Articles 15–21), and the right to lodge a complaint with a supervisory authority. Where Türkiye's KVKK (Law No. 6698) applies, you additionally have the rights in its Article 11 — including learning whether your data is processed, requesting correction or deletion, and objecting to results produced exclusively by automated analysis. Much of this is self-service:

- Access & portability — the full account export (ledger, job history, custody records, lineage) is available from the dashboard at any time, including while an account is frozen for deletion.
- Erasure — per-project retention knobs for content; the account-deletion flow (section 7) for everything.
- Rectification, restriction, objection — contact us (section 14) for account data; for customer content, requests go to the controller (you or the organisation that submitted it), and we assist under the DPA.

If we receive a data-subject request concerning content processed for one of our customers, we will redirect the requester to that customer and assist the customer as the DPA describes.

## 11. Who we share with

We share personal data only with the service providers on our subprocessor list, each under terms consistent with the DPA; with the generation providers and storage vendors you connect (your instruction, your contract — they are not our subprocessors); and where the law requires disclosure. We do not sell personal data.

## 12. Security

Security measures are described on the Security & Trust page: per-secret envelope encryption of credentials under a managed key-management service with tenant-bound encryption context, tenant isolation enforced in the data layer, an append-only audit ledger, byte-exact sha256-verified custody, private buckets with short-lived signed URLs, and secrets kept out of logs. Wakemark holds no security certification today (SOC 2 is on the roadmap, not a current attestation) and claims none.

## 13. Children

Wakemark is a business tool for professional use, not directed to children, and we do not knowingly collect children's personal data. The Terms of Service set a minimum age for account holders.

## 14. Changes to this policy

The last-updated date at the top reflects the most recent change. We will give notice of material changes before they take effect — on this page and, for account holders, by email.

## 15. Contact

Reach us at hello@wakemark.ai. The controller for account data is O ÷W¦† â ±1r, an individual (a natural person) based in Türkiye. Data-protection and privacy inquiries — including the exercise of the rights in section 10 — are handled at the same address. Where an EU representative (GDPR Article 27) is required, we will designate one and name it here.