Taproot

Subprocessor List

Version 1
Effective November 13, 2026

Taproot IO, LLC uses the following subprocessors to operate the service. Vendor contracts and production accounts are authoritative; this list is reconciled against them before each published version.

DigitalOcean, LLC

Purpose: Cloud infrastructure: Kubernetes compute, managed PostgreSQL database, Spaces object storage (the origin for published sites and images), container registry, load balancing, networking and backups

Data processed: All service data: account, content, membership, billing-reference, operational, and log data

Location: United States (the cluster, database, and Spaces are in the San Francisco region, sfo2); support access per DigitalOcean terms

Cloudflare, Inc.

Purpose: DNS, proxy, CDN and edge delivery of the Taproot app, published sites and images; edge compute and operational request logs (Workers and Workers Logs); routing data (Workers KV and a Durable Object); analytics measurement store (Analytics Engine); creator-uploaded video storage and delivery (R2 and a cached video host); custom-domain TLS issuance and per-hostname routing (Cloudflare for SaaS and API); DNS verification lookups

Data processed: Every request that crosses the proxy (IP address, user agent, headers), including the content of the visitor calls a published site sends to Taproot's API (form submissions, comments, member sign-in, checkout). Cached published content; site and release routing records; creator-uploaded video files. Stored request logs can include URLs and query strings, request headers containing visitor IP addresses, referrer, user agent, request identifiers, network and location metadata. No blanket header, query-string or cookie redaction is asserted. Workers Logs keeps request logs for up to seven days, separately from Analytics Engine and Honeycomb. Analytics events keyed by a pseudonymous identifier derived from request data (not the raw IP address or user agent), with page, previous-page, referrer, country and device fields, and, where the site's analytics features are on, a monthly reader identifier, visit-context fields and engagement measurements; summaries are kept as counts. The provider keeps individual analytics events for up to three months

Location: Global network. Taproot has not restricted where Cloudflare stores data: Workers KV is replicated globally, the Durable Object and Analytics Engine have no region set, and the video bucket is created with a Western North America location hint and no jurisdictional restriction

Amazon Web Services, Inc. (SES/SNS)

Purpose: Transactional and creator email delivery and delivery/bounce/complaint event notifications

Data processed: Recipient and sender addresses, message content, delivery and suppression metadata

Location: United States (us-east-1)

Stripe, Inc.

Purpose: Platform subscriptions, member payments, tax calculation, checkout, Connect Express onboarding and payouts, fraud prevention

Data processed: Identity, contact, billing, transaction, tax, connected-account, fraud, and dispute data

Location: United States; Stripe group entities per Stripe's terms

Honeycomb.io, Inc.

Purpose: Operational monitoring: warning-level and error logs, sampled traces, Kubernetes warning events, and system metrics from production; sanitized delivery-error events from the edge workers

Data processed: Pseudonymous user, account, site, and request identifiers; error text with email addresses, tokens, and web addresses removed before export; for edge events, a site identifier, cache result, status, and a hash of the host and path, with no IP address or user agent. Events and logs are kept 60 days, metrics 13 months

Location: United States

Google LLC (Places API)

Purpose: Location search and place details for place-related content authoring

Data processed: User-typed location search queries, plus an approximate (city-level) location derived from the creator's IP address by Cloudflare and used only as a result-ranking bias, not stored; proxied through Taproot's servers without account identifiers. The place a creator selects (name, address, coordinates) is stored with their content

Location: United States/global

Google LLC (Search Console API and OAuth authorization)

Purpose: Optional: when a site's owner connects Search Console, Taproot imports that site's search performance data

Data processed: The owner's Google authorization tokens (stored encrypted) and the verified site address; Google returns search queries, pages, clicks, impressions, and positions for the site. Daily data is kept 90 days and monthly data 16 months

Location: United States/global

Google LLC (Business Profile APIs and OAuth authorization)

Purpose: Optional: when a site's owner connects a Google Business Profile, Taproot reads that location's details and reviews, and posts a page or a review reply to the profile when a site manager confirms one

Data processed: The owner's Google authorization token (stored encrypted). Google returns the locations the signed-in Google account manages, each with its name, address, phone number, website, Google's place identifier and its maps and review links; the list is held for up to about an hour while the owner chooses. For the chosen location Google also returns its reviews: the text, star rating, dates, the owner's reply and, when the reviewer is not anonymous, the reviewer's display name and profile photo. Only the chosen location's details and reviews are kept, while the connection lasts, and they are removed after 21 days without a refresh. Taproot sends Google the posts and review replies a site manager confirms: the text, an optional image and a link to the page

Location: United States/global

Hyperknot Software Kft. (OpenFreeMap)

Purpose: Map tiles for a published site's map page, requested only when a visitor opens that page

Data processed: The visitor's IP address, browser information and the map area viewed

Location: Hungary (operator); tiles may be delivered through Cloudflare's global network

Proton AG (mail)

Purpose: Receives mail sent to Taproot's own addresses (privacy@ , legal@ , support@ , and copyright@taproot.io)

Data processed: The sender's address and whatever the message contains, including data-subject requests and copyright notices

Location: Per Proton's terms

Image processing (imgproxy) runs on Taproot's own infrastructure inside DigitalOcean and is not a separate subprocessor. GitHub, Inc. hosts source code and runs the build/deploy pipeline; it does not process user personal data. Building the service also downloads software and font files from public sources such as jsDelivr, Docker Hub, GitHub Container Registry, and Google Fonts; those requests come from Taproot's servers and carry no personal data, and published pages never load fonts or scripts from them.

User-directed and independent services

Services a creator or user independently chooses — creator-linked destinations, embedded third-party content such as a YouTube or Vimeo video (loaded only after a reader presses play), integrations a creator installs, domain registrars used to configure DNS for a custom domain — act under their own terms rather than as Taproot subprocessors. Taproot has no third-party sign-in providers; authentication is first-party.

The bot challenge on account sign-up and published-site forms is self-hosted (Cap, running in Taproot's own cluster), so it is not a subprocessor: a visitor's browser talks only to Taproot's own origin, no challenge vendor receives any data, and the challenge service keeps only short-lived challenge state and a rate-limit counter keyed by the visitor's address for a few seconds. Taproot's own database also keeps a ten-minute counter keyed by a hash of the visitor's address to limit repeated submissions, and a form submission keeps that hash for 30 days (see the Privacy Notice).

Change process

Before adding a provider that will process personal data:

1. Complete security, privacy, location, retention, and contract review.
2. Record service purpose, data categories, legal role, regions, and transfer mechanism in the controlled vendor inventory.
3. Determine whether customer notice or an objection window is required (DPA customers receive at least 30 days' notice of material additions or replacements).
4. Update and deliberately publish a new version of this list.

Questions about this list: privacy@taproot.io . Taproot publishes a new dated version of this list before a new provider begins to process personal data, and gives customers who have signed the Data Processing Addendum the notice that addendum requires.