SendTax

TRUST CENTER

Data Retention & Deletion Policy

How long SendTax retains customer data, what triggers deletion, and how to exercise access and erasure rights.

Effective: May 21, 2026 Owner: Howell & Gibbs LLC Status: v1 – initial publication

Review cadence: Semi-annual, or after any material change to data categories, regulatory obligations, or disposal mechanisms

---

1. Purpose

This policy defines how long SendTax retains the data it holds, what triggers deletion, and how customers can exercise their rights to access, correct, and erase their information. It supports the commitments made in the **Information Security Policy** and aligns with SOC 2 Common Criteria CC6.5, C1.1, C1.2, and Privacy criterion P4.2.

SendTax handles taxpayer data subject to specific regulatory retention obligations. This policy reflects those obligations and adds SendTax-specific commitments where appropriate. It is a companion to SendTax's Privacy Policy (which governs *what* data is collected and *why*) and to the Information Security, Access Control, Encryption, and Incident Response policies.

2. Principles

- **Purpose-bound retention.** SendTax retains personal data only as long as there is a defined purpose for holding it — service delivery, legal compliance, or product-improvement work the customer has explicitly consented to.
- **Legal obligations override convenience.** Where U.S. or other applicable tax-data law requires retention longer than the customer might prefer, the legal minimum prevails until the obligation expires.

- **Right to know, right to correct, right to delete.** Customers may request access, correction, and (subject to §6) deletion of their data, on the timelines required by applicable law.
- **Documented purges.** Automated retention sweeps run on a defined cadence, log their actions, and are auditable.
- **Honest disposal.** When SendTax says data is deleted, the disposal method makes that true, with the exception of backup-aging behavior described in §7.
- **Customer control where feasible.** Customers can delete individual documents and request account deletion through the mechanisms in §4.

3. Categories of data and retention periods

SendTax categorizes data by purpose. Each category has a retention period set by a combination of business need, legal obligation, and customer consent. Periods marked **(roadmap)** are commitments SendTax is implementing; current behavior is described in §4 and §9.

3.1 Service data (the default category)

What it is: account profile, contact details, tax-return data being prepared, supporting documents uploaded by filers, audit-log entries for activity on the account.

Retention: 7 years after the last activity on the account, unless a longer retention period is required by law or contract. *(Roadmap — automated sweep implemented; runs in dry-run pending production enablement; see §9.)*

Why this number: Matches AICPA professional guidance for tax-data retention; exceeds the IRC §6107(b) three-year minimum applicable to certain preparer records; accommodates most state tax agency requirements and typical IRS statute-of-limitations windows. Customers may request earlier deletion under §5, subject to the legal-obligation carve-outs in §3.2 and the limits described in §5.

3.2 IRS-mandated e-file records

What it is: taxpayer authorizations and supporting records that the IRS requires authorized e-file providers to retain — including Form 8878 (e-file signature authorization for application for extension), Form 8879 (e-file signature authorization for individual returns), and any equivalent state authorizations.

Retention: 3 years from the later of (a) the return due date or (b) the IRS-received date, per IRS Publication 1345.

Override: This retention obligation overrides any customer request to delete the affected records earlier. SendTax will identify the specific records affected, the statute, and the date the obligation expires, and will honor deletion of any records not covered by the obligation.

3.3 Identity documents

What it is: documents used for identity verification (e.g., driver's license images, identity verification scans).

Retention: Retained only as long as needed for identity verification; target deletion is verification completion + 90 days. *(Roadmap — identity purge implemented; runs daily in dry-run pending production enablement; see §9.)*

Why this number: Data minimization. Identity documents have a narrow operational purpose and should not be retained long-term.

3.4 Anonymized training contributions

What it is: anonymized document artifacts that customers explicitly contributed to the SendTax model-improvement program, stored in Cloudflare R2 under the contributions bucket prefix.

Retention: 1,095 days (3 years) from contribution approval, configured by the `TRAINING_ARTIFACT_RETENTION_DAYS` setting and enforced by an automated daily purge task (`purge-expired-training-contributions`). At expiration, the R2 artifact is deleted, the wrapped DEK and storage pointer are cleared on the contribution row, a `purged_at` timestamp is stamped, and the labeler image cache is invalidated.

Important distinction: Derivative label metadata (annotations describing what was in the document) is retained for model-lineage purposes; the underlying document is gone. SendTax does not retain identifying information in the label metadata.

Why this number: Aligned with FTC Safeguards Rule retention-minimization expectations and with the consent text customers see when they opt in.

Customer override: Contributors may withdraw a contribution at any time through the separate withdrawal flow, which deletes the artifact and removes the labels regardless of the 3-year window.

3.5 Authentication and identity records

What it is: account identifiers, sign-in history, MFA enrollment state. Most of this data is held by Clerk (SendTax's identity provider) rather than in SendTax-controlled databases.

Retention: Governed by Clerk's data-handling terms while the SendTax account is active. On account closure, the Clerk-side identity is removed via the Clerk API; the SendTax-side mirror record is soft-deleted immediately (see §4) and hard-deleted at the expiration of the §3.1 window.

3.6 Application audit logs

What it is: server-side records of security-relevant actions — sign-in events, role changes, administrative actions, sensitive data access. See **Information Security Policy** §8.

Retention: Retained for at least **7 years**, matching the longest tax-data retention obligation SendTax is subject to, so audit trails covering retained records always outlive those records. Audit logs may be retained longer for ongoing investigations or litigation holds.

3.7 Incident response records

What it is: declared security incidents and their append-only update timelines, as described in the **Incident Response Policy**.

Retention: Retained **indefinitely**. The audit value of a complete incident history outweighs the storage cost, and customers benefiting from past notifications should always be able to see the public record of what happened.

3.8 Operational telemetry

What it is: error reports (Sentry), application performance metrics, request logs.

Retention: Per the provider's default retention. Scrubbed of known sensitive fields before transmission. **Sentry events:** retention is governed by the active Sentry plan and per-project settings; the working assumption is **90 days** consistent with the Team plan default, to be confirmed and pinned at publication. **Fly.io logs:** retained per Fly's published platform behavior (approximately 3 days of live log retrieval via `fly logs`); longer retention is not currently configured and would require a log shipper. *(2026-05-26 status: Liam to confirm the exact Sentry org-settings retention and Fly platform retention in their respective dashboards, then update this section to the actual numbers before publication.)*

3.9 Customer support communications

What it is: support emails and in-product feedback messages (FAB tickets).

Retention: 3 years.

3.10 Billing records

What it is: invoices, payment records, and tax records associated with SendTax's own business operations.

Retention: 7 years, consistent with SendTax's own tax-compliance obligations.

3.11 Marketing and website analytics

What it is: non-identifying website analytics (e.g., Google Analytics 4 page views).

Retention: 14 months, matching Google Analytics 4 default.

3.12 Backups

What it is: point-in-time database snapshots produced by Fly.io Managed Postgres, and R2 object durability.

Retention: Per Fly.io Managed Postgres's published retention window; SendTax does not maintain a separate backup pipeline. Backups inherit any retention obligation of the underlying record but, because they are point-in-time and not per-record, SendTax does not selectively expunge an individual customer's data from historical backups. Backups age out of the retention window on the provider's schedule. See §7 for the practical implications of backup persistence on customer deletion requests.

4. The deletion lifecycle

When a customer deletes their SendTax account, deletion happens in two stages.

4.1 Stage 1 — Immediate (soft delete)

On receipt of the `user.deleted` event from Clerk (which fires when a customer deletes their account through Clerk's hosted account UI), or on equivalent customer-initiated action:

- The Clerk-side identity is removed by Clerk.
- The SendTax-side `User` or `Filer` mirror record is marked `is_active = false`.
- All active sessions for the principal are revoked.
- The account becomes inaccessible to its former owner; existing firm-to-filer links are moved to the `ENDED` state.

At this point the customer's data is no longer visible or usable in the product. The retention period for any records subject to a legal obligation (e.g., IRS-mandated e-file records) continues to run.

4.2 Stage 2 — Permanent (hard delete)

After the applicable retention period elapses:

- Tax-return data and supporting documents no longer subject to a legal retention obligation are hard-deleted.
- The R2 storage objects for those documents are deleted.
- The wrapped DEKs for those documents are removed along with their database rows, rendering the document ciphertext unrecoverable from SendTax systems (see §8).
- Audit-log entries describing those records are retained for the separate audit-log retention window (see §3.6) but reference only opaque identifiers, not document contents.

The automated sweep that performs Stage 2 is **implemented** and runs daily on the worker schedule in **dry-run mode** — logging the accounts, documents, and identity records it would purge, with a queryable audit row per candidate — pending production enablement of its retention flags (`RETENTION_HARD_DELETE_ENABLED`,

`RETENTION_DOCUMENT_PURGE_ENABLED`, `IDENTITY_DOCUMENT_PURGE_ENABLED`). Customer-requested deletions are honored on the §5 timelines regardless of the sweep's mode.

4.3 Individual document deletion

Filers can delete individual documents from their account at any time. When a filer deletes a document:

- The document's database row is removed, which removes the wrapped DEK stored on that row.
- The corresponding object in Cloudflare R2 is deleted.
- The action is recorded in the audit log.
- Because the wrapped DEK is unique to the document and is removed along with the row, the underlying ciphertext is rendered unrecoverable from SendTax systems, independent of when R2 finishes propagating the object deletion.

4.4 Cascade behavior

When a database record that owns dependent records is deleted, SQLAlchemy `cascade="all, delete-orphan"` ensures dependent rows are deleted with the parent. This applies to relationships including User → Filer, Filer → Document, and User → External Identity. Cascading deletion is the standard data-integrity behavior; it does not by itself constitute disposal under this policy unless the parent record is itself being disposed of.

5. Customer rights

SendTax honors the following data-subject rights, consistent with applicable privacy law (including GDPR, UK GDPR, CCPA/CPRA, VCDPA, CPA, CTDPA, UCPA, and other comparable statutes where applicable):

Right	How to exercise it	Timeline
Access — receive a copy of personal data SendTax holds	Email privacy@send.tax from the account's registered	Acknowledged within 5 business days; substantive response within 30 days

	address, or use in-product export tools where available	
Correction — fix inaccurate data	Self-service in product where possible; otherwise privacy@send.tax	Within 30 days
Deletion / erasure — remove personal data	Account deletion in Clerk's hosted account UI, or written request to privacy@send.tax	Stage 1 immediate; Stage 2 at retention-window expiry
Restriction / objection — limit how personal data is processed	privacy@send.tax	Within 30 days
Data portability — receive personal data in a machine-readable form	privacy@send.tax	Within 30 days
Withdraw consent — for processing based on consent (e.g., training contributions)	In-product withdrawal flow, or privacy@send.tax	Immediate

5.1 Limits

Some deletion requests cannot be honored in full while a legal obligation applies. Where that is the case, SendTax will:

1. Acknowledge the request within the statutory timeline.
2. Identify the specific records covered by the obligation.
3. Identify the specific statute and the date the obligation expires.
4. Honor deletion of any records not covered by the obligation.
5. Automatically complete deletion of the held-back records when the obligation lapses.

6. Verification of deletion requests

To prevent unauthorized erasure, SendTax requires deletion requests to be verifiable. Standard verification methods:

- A request sent from the email address associated with the SendTax account
- A request made while signed in to the SendTax product
- For requests received through other channels, a brief challenge-response confirming control of the account's contact details

SendTax does not require government-issued identification for routine deletion requests, in keeping with data-minimization principles. Additional verification may be required for requests touching IRS-mandated records or unusually large data exports.

7. Backups

Customer data deleted from SendTax production systems may persist in Fly.io Managed Postgres backups until those backups are naturally aged out per the managed service's retention schedule. SendTax does not selectively edit or scrub backups, as this is operationally infeasible and not industry-standard practice.

The practical implication: when a customer deletes their data, the data is removed from active systems immediately, but a copy may exist in backup storage for the duration of Fly.io's backup window. After the backup ages out, no recoverable copy exists on SendTax-controlled systems.

Cloudflare R2 object deletion takes effect according to R2's published behavior. SendTax does not maintain a separate backup of R2 objects; document deletion in the application immediately deletes the R2 object.

8. Disposal methods

SendTax uses the following disposal methods, applied as appropriate to the data category:

- **Database row deletion.** A `DELETE` statement removes the row and any cascaded child rows. Used for hard-deletion of customer data, account records (post-grace), audit logs past retention, and other structured records.

- **Per-document DEK invalidation through row deletion.** When a document database row is deleted, the wrapped DEK stored on that row is destroyed with it. Because each document has a unique DEK that exists nowhere else, this renders the document's ciphertext unrecoverable from SendTax systems, regardless of whether the underlying R2 object has been fully propagated. SendTax does not currently destroy Key Encryption Key (KEK) versions in Google Cloud KMS as a routine disposal mechanism, since per-document DEK uniqueness makes this unnecessary for document-level deletion. KMS key-version destruction remains available for bulk-deletion scenarios that may arise (for example, mass tenant offboarding) and would render all documents wrapped under that KEK version simultaneously unrecoverable.
- **Object storage deletion.** R2 objects are deleted via the R2 API. Cryptographic protection (§4.3) provides defense against any propagation lag.
- **Soft-delete flagging.** Records that need a grace period (e.g., user accounts) are flagged inactive in the database while retaining the underlying data, until the grace period passes and the record is eligible for hard-deletion.
- **Provider-side aging.** For data held by sub-processors (Clerk auth logs, Postmark email records, Sentry telemetry, Fly.io backups), disposal occurs per the provider's published retention behavior.

9. Automated retention enforcement

Retention is enforced by code, not by memory. SendTax operates the following automated mechanisms; those marked *dry-run* are built and on the daily schedule but gated off pending production enablement:

Mechanism	Cadence	Status	Source of truth
Training-contribution artifact purge	Daily	✅ Production	`purge-expired-training-contributions` Celery Beat task in `st-backend/app/workers/training_purge_tasks.py`
Soft-delete on account deletion	Real-time	✅ Production	Clerk `user.deleted` webhook in `st-backend/app/api/routes/webhooks_clerk.py`

Session revocation on account deletion	Real-time	✅ Production	Same webhook handler
Hard-delete sweep of soft-deleted accounts past retention window	Daily	✅ Built · dry-run	Implemented and on the daily beat in dry-run; enable via `RETENTION_HARD_DELETE_ENABLED`. `purge-soft-deleted-accounts` in `st-backend/app/workers/account_retention_tasks.py`. Production-enablement target: 2026-07-26.
Document retention sweep (7-year purge of inactive accounts)	Daily	✅ Built · dry-run	Implemented and on the daily beat in dry-run; enable via `RETENTION_DOCUMENT_PURGE_ENABLED`. `purge-expired-documents` in `st-backend/app/workers/document_retention_tasks.py`. Production-enablement target: 2026-08-26.
Identity-document deletion after verification + 90 days	Daily	✅ Built · dry-run	Implemented and on the daily beat in dry-run; enable via `IDENTITY_DOCUMENT_PURGE_ENABLED`. `purge-expired-identity-documents` in `st-backend/app/workers/identity_retention_tasks.py`. Production-enablement target: 2026-07-26.
Advance notice to customers before permanent deletion	Daily	✅ Built · dry-run	Implemented and on the daily beat in dry-run; enable via `RETENTION_NOTICES_ENABLED`. `send-retention-notices` in `st-backend/app/workers/retention_notice_tasks.py`. Production-enablement target: 2026-08-26.

Every purge run logs structured events recording what was deleted and how many records were affected, enabling after-the-fact audit. Once each dry-run mechanism is enabled in production, this policy will be revised to reflect active enforcement and the corresponding marker removed.

10. Special situations

10.1 Litigation hold

If SendTax receives a credible litigation hold, governmental preservation order, subpoena, or comparable legal instruction, the affected records are placed on hold and exempted from automated retention sweeps for the duration of the hold. The Communications Lead (or Security Lead, where applicable) is responsible for:

- Identifying the records covered by the hold
- Notifying anyone with disposal authority over those records
- Documenting the hold's scope, basis, and expected duration
- Releasing the hold and resuming normal retention once the legal basis ends
- Reviewing the hold at least quarterly to confirm it remains in effect

Records under legal hold are not deleted even if they have aged past their normal retention period.

10.2 Sub-processor data

Personal data that flows to a sub-processor (e.g., a document held in R2, an email sent via Postmark) is governed by both this policy and the sub-processor's own retention terms. Where customer-initiated deletion at SendTax results in corresponding deletion at a sub-processor, the cascade is engineered into the integration. Sub-processor copies that age out per the provider's own retention schedule are out of SendTax's direct control. The current list of sub-processors is published on SendTax's Sub-Processor List.

10.3 Anonymized and aggregated data

Once data has been irreversibly anonymized — meaning it can no longer be associated with a specific individual even when combined with other data SendTax holds — it is no longer "personal data" for the purposes of this policy and may be

retained indefinitely. Examples: fully de-identified statistics about product usage, model-evaluation metrics derived from training contributions, error counts.

11. Exceptions

Any deviation from this policy — for example, extending retention of a specific record set for an active investigation or compliance review — requires:

1. Approval from a second SendTax operator
2. A documented justification recorded with the change
3. A defined re-evaluation point

Exceptions are reviewed at the next policy review and are not allowed to become permanent without explicit re-authorization.

12. Enforcement

Violations of this policy may result in revocation of access, termination of employment or contract, and, where applicable, civil or criminal referral.

13. Related policies

- **Information Security Policy** — umbrella policy
- **Access Control Policy** — audit log retention floor (§8 of that policy) and customer account deletion (§4.6)
- **Encryption Policy** — disposal mechanisms involving wrapped DEKs (§8.5)
- **Incident Response Policy** — incident record retention (§11)
- **Privacy Policy** — customer-facing rights, the legal basis for processing, and the consumer-facing summary of this retention schedule

14. Document control

Version	Date	Author	Notes
1.0	May 21, 2026	Holly Gibbs	Initial publication

15. Contact

Privacy / data subject requests	privacy@send.tax
Security disclosures	security@send.tax
Operating entity	Howell & Gibbs LLC (operator of SendTax)