

Privacy Policy

What personal data Tephlo handles, why, for how long, who else sees it, and what you can ask us to do about it. Written to match what the software actually does.

Last updated 6 September 2026
Revision Draft 1.1 — factual content refresh
Version 2 · Effective 6 September 2026

TEPHLO AI ENTERPRISES

9A Kenneth Lekara Luka Street, Lokogoma, FCT, Nigeria

Tephlo · hello@tephlo.com

<https://legal.tephlo.com>

Exported from the published pages on 2026-10-02T17:17:50.686Z. The pages are the authoritative text.

IN PLAIN LANGUAGE

Tephlo is used by businesses to answer their customers. That means two very different sets of people, with two different answers.

- **If you messaged a business** that uses Tephlo: your messages and the phone number, Telegram id or random web id you wrote from are stored so the business can answer you and keep the history. The business decides why — we handle it on their instructions, and they can export or delete your data.
- **If you run a workspace:** we hold your name, work email, a hashed password, your role and sign-in records, because that is what an account is.
- Messages are sent to an AI provider to generate a reply. They are not used to train AI models, and they are never sold.
- Some conversation context lives only in short-lived memory — the in-conversation working state expires after 30 minutes.
- This site sets no tracking cookies at all. See the [cookie notice](#).

1. Two different roles, and why it matters

TEPHLO AI ENTERPRISES operates Tephlo. Depending on whose data is involved, we act in one of two roles:

Processor — for the conversations businesses have with their customers

A business (our customer, called a “workspace” here) decides to run an assistant, decides what knowledge it answers from, and decides how long conversations are kept. They are the controller. We store and process that data on their instructions, under the [Data Processing Agreement](#). If you are a customer of that business and want your data removed, the fastest route is to ask them — they hold the controls.

Controller — for the accounts of the people who run a workspace

When someone signs up, is invited to a workspace, signs in, contacts support, or visits our public sites, we decide why that data is handled. That is what sections 4, 5 and 6 of this policy describe.

2. What is collected from people who message a business

A workspace can answer on WhatsApp, Telegram, a website chat widget or an enabled commercial API integration. Depending on the features used, the platform holds:

- **Your channel identifier.** On WhatsApp this is the phone number the message came from; on Telegram, your Telegram user id; in the web widget, a random identifier the widget generates in your browser and reuses on that site. An API integration supplies its own customer identifier. These identifiers link messages to a conversation; a random widget id does not require a name, email or account.
- **Message content.** What you wrote and what was sent back, including replies written by a human agent, with timestamps and the provider's message id (used to avoid processing the same message twice).

- **Conversation metadata.** The channel, which business identity received the message, whether the assistant or a person is handling it, when it was last active, and — where a conversation was escalated — the reason and whether it was resolved.
- **Working state during a conversation.** What you appear to be trying to do, the option you chose, details you already gave, and the question last asked, so nobody has to repeat themselves. This is held in fast temporary storage and expires 30 minutes after the last message.
- **Routing records, on shared numbers.** Where a business is reached through a number the platform shares between businesses, a record of which business a conversation was routed to is kept using a salted one-way hash of your identifier rather than the identifier itself.
- **Captured requests.** If you ask for a callback, make a complaint or request a refund, a ticket is created for the team with a short summary of what you asked for.
- **Remembered facts, where customer memory is allowed.** Short statements you make about yourself — a name, a preference, a product you own — can be stored so you do not have to repeat them next time. Customer memory is enabled by default in the software; the workspace can switch it off, and sensitivity controls can block it. The current extraction path accepts stated facts, filters instruction-like and secret-like content, and uses at most twelve facts in a reply. The assistant can still misunderstand a statement; administrators can inspect and forget stored facts.
- **Media you send, where those features are enabled.** Voice notes can be transcribed to text; images can be interpreted; documents can be scanned for malware and read for the answer. Where spoken replies are enabled, the approved reply text is sent to a speech provider to produce audio. These optional features require the relevant platform, provider and workspace settings. What a shared location contains is used at the moment it arrives to work out which of the business's branches is nearest; a shared contact card is acknowledged and is never saved to an address book or contacted automatically.
- **Carts, requests and appointments.** Saved carts contain chosen products, quantities and price snapshots. Potential-order requests record a request for workspace review, not a placed or paid order. Appointment records contain the selected service, resource, time and request or booking status. Structured care cases can hold a complaint, the requested outcome and the team's decisions.
- **Smart Updates preferences and history.** Where enabled, the platform records the topic you agreed to receive, the consent text and time, preferences, opt-outs and updates sent. A conversation alone is not consent to updates. You can reply STOP to stop them.
- **Operational measurements.** Counts and categories — how a turn ended, how many tokens a reply used, how long it took. These carry no message content. Some records are linked to a conversation so they can be removed with it; this does not make them anonymous.

3. What the assistant refuses to collect

Some data is refused in code rather than discouraged in policy. The assistant will not ask you for, and will not store as conversation state, fields it recognises as payment instruments, credentials, bank details or government identity numbers — card and CVV numbers, PINs, passwords, one-time codes, IBAN and account numbers, passport and national identity numbers, dates of birth. The rule recognises these in English, French and Spanish, because the platform's first market is francophone and a filter that only speaks English only protects English speakers.

Where text is stored for optional quality review, or extracted from a document you send, recognisable emails, phone numbers, long card-like numbers and labelled codes are replaced with markers before it is stored.

Being exact about this. Pattern matching is a strong filter, not a guarantee. If you type a card number into a chat unprompted it may be stored as ordinary message text. Never send card details, passwords or one-time codes in a chat message — no legitimate business needs them that way.

4. What is collected from the people who run a workspace

- **Account details.** Name, work email address, a hashed password, role, whether the account is active, and — if you choose to use them for password recovery — a phone number or linked Telegram chat.
- **Sign-in and security records.** Session and refresh-token records, multi-factor authentication state (the secret itself is stored encrypted), failed sign-in counters and lockouts.
- **Signup details.** For self-serve signup, the email address, company name and display name given, held against a verification link until it is used or expires.
- **Administrative audit records.** Who changed a channel credential, who ran a privacy export or deletion and the reason they gave, who changed routing or platform configuration. These record the action, not the content.
- **Email delivery records.** Verification, invitation, password reset and escalation-alert emails are recorded with the recipient and subject; the body is stored encrypted and cleared once delivery reaches a final state. Escalation alerts deliberately carry no customer message content — they say a conversation needs a person and link to the console.
- **Workspace billing records.** Where plan billing is enabled, payer email, selected plan, amount and currency, payment references, status and renewal records are retained. A provider may return a reusable authorization reference and limited display details such as card brand, last four digits and expiry. Full card numbers are entered on the provider's hosted checkout, not collected by the platform's checkout flow.
- **Support correspondence.** What you send us when you ask for help.

5. Website visitors

Our public sites — the marketing site, the documentation site and this legal site — run no analytics, no advertising scripts and no third-party tracking of any kind. Their content security policy blocks third-party scripts outright. The only browser storage involved is a currency preference you set yourself on the pricing page. The [cookie notice](#) lists everything, including the session cookies the signed-in consoles need.

Our hosting and infrastructure providers process connection data such as IP addresses in server logs for security and abuse prevention, as any web host does.

Published logos and icons can be served from configured asset hosts, which receive ordinary connection data. If you submit the contact or demo form, its contents are sent to the backend so the enquiry can be recorded and answered; viewing the form does not submit it.

6. Why we handle this data, and on what basis

For customer conversations, the workspace you contacted decides the purpose and the lawful basis. We process on their documented instructions and tell them what the platform makes possible.

Where we are the controller, our purposes and bases are:

PROCESSING WHERE WE ARE THE CONTROLLER

Purpose	Data	Basis
Providing accounts and the consoles	Account details, sign-in records	Performance of our contract with your business
Securing the platform	Authentication records, audit trails, lockout counters, error reports	Legitimate interests — keeping a multi-tenant platform safe
Operational email	Email address, delivery records	Performance of the contract; consent for escalation alerts, which are confirmed by email first and can be stopped at any time
Support and correspondence	What you send us	Legitimate interests — answering you
Meeting legal obligations	Records we are required to keep	Legal obligation

We do not use personal data for advertising, we do not sell it, and we do not use customer conversations to train AI models.

7. How long data is kept

These are software defaults, not a statement that every deployment has its retention workers enabled. Read-time expiry stops expired data being used; physical removal by a scheduled sweep requires that the responsible worker is running. Available controls and periods vary by record type. An explicit customer erasure is a separate operation.

What	Kept for	Then
In-conversation working state	30 minutes after the last message	Expires automatically; deleted immediately on erasure
Web chat widget session	1 hour of inactivity	Expires; the widget starts a new session
Conversations the assistant handled	365 days after the last activity, by default	Deleted by a scheduled sweep, with their messages and escalations Conversations waiting for a person, or being handled by one, are not age-deleted — they stay until resolved or explicitly deleted.
Delivery records for outbound messages	7 days after publication, by default	Deleted

Duplicate-delivery protection records	30 days after completion, by default	Deleted
Customer documents processed for text extraction	Deleted as soon as the text has been extracted, by default	Extracted text is stored encrypted for 7 days by default and can be configured to never be stored at all
Optional quality-review excerpts	30 days by default, where a workspace enables them	Deleted; sensitive values are masked before storage
Remembered customer facts	Newly extracted facts: 180 days; administrators can forget them or erase the customer earlier	Expired facts are excluded when memory is read. Physical deletion depends on the retention worker running. Legacy entries without an expiry remain until explicitly forgotten or erased. Replacing a stated fact can refresh its expiry.
Staff accounts	For as long as the account exists	Deleted with the account or the workspace
Administrative audit records	Kept as platform evidence	Survive workspace deletion; contain no message content
Encrypted database backups	Rolling schedule — hourly, daily and monthly copies	Pruned as they age out; deleted data leaves backups then, not at the moment of deletion

8. Who else receives this data

The [sub-processor page](#) is the inventory of provider roles, with the actual recipients, regions and contractual details to confirm for the deployment. In outline:

- **An AI provider** receives the conversation and the relevant passages of the workspace's knowledge in order to generate a reply.
- **WhatsApp (Meta) and Telegram** carry the messages themselves — they see everything sent over their own channel, under their own terms.
- **Hosting, database and storage providers** hold the data at rest.
- **An email provider** delivers verification, invitation, password reset and alert emails to staff addresses.
- **A billing provider**, where configured, handles hosted payments for workspace plans and receives payer and transaction details. This is separate from a customer submitting a potential-order request.
- **Optional providers** — transcription, spoken replies, image reading, web search and error monitoring — are off unless configured and are listed with their status.

We may also disclose data where the law requires it, or to establish or defend legal claims. If our business is transferred, data moves with it and this policy continues to apply until it is replaced.

9. International transfers

The providers above operate internationally, and the AI provider may route a request to a model provider in another country. We cannot promise that data stays inside one country or region, and this policy will not pretend otherwise.

Where personal data protected by the law of the Nigeria Data Protection Act 2023 leaves that jurisdiction, the transfer mechanism relied on is **[TRANSFER MECHANISM]** — this must be confirmed by counsel for each provider before this policy is published as final.

10. Your rights, and how they work here

Depending on where you are, you may have rights to access your data, correct it, delete it, restrict or object to its use, receive a portable copy, and withdraw consent you gave.

If you messaged a business that uses Tephlo, ask that business: they are the controller, and their administrators have two tools built into the console.

- **Export** produces a structured copy of one customer's data with that business: every conversation, the full transcript, escalations, shared-number sessions, routing records, memberships, remembered facts, tickets, care cases, customer-scoped cached answers, Smart Updates consent and send history, carts, potential-order requests and appointments. Review excerpts are exported as metadata; their message text is already in the transcript.
- **Deletion** permanently removes conversations and their messages, escalations, transcripts of media in them, shared-number sessions, routing records, memberships, queued outbound messages and in-flight processing records, and clears the short-lived working state. Work already accepted before the deletion is suppressed rather than delivered afterwards, so a reply cannot arrive after someone asked to be forgotten. Both actions are recorded in an audit trail that identifies the subject only by a salted hash.

What deletion covers. One request removes the conversations and their messages, the escalations raised from them, the routing and session records, any queued deliveries, the facts remembered about that person by the customer-memory feature, tickets and care cases, customer-scoped cached answers and review excerpts, Smart Updates consent and send history, saved carts, potential-order requests and appointments. Related operational records and withheld reply drafts are erased too; they are not all separately disclosed in an export. What remains includes an audit record of the deletion itself, which identifies the person only by a salted hash and exists so the erasure can be proven to have happened.

If you run a workspace, write to us using the contact details below. We will verify who you are before acting — usually by using an address already on the account — and answer within one month, or tell you why we need longer.

11. How the data is protected

The [Security Overview](#) is the detailed answer. The short version: each workspace's data is separated and every request is scoped to the account making it; connections are encrypted; channel credentials, authentication secrets, email bodies and extracted document text are encrypted at rest; passwords

are hashed; sensitive values are stripped from logs; administrative actions are recorded in audit trails, two of which the database itself refuses to let anyone alter.

What we do not claim. There is no SOC 2 report, no ISO 27001 certificate, and message content is not end-to-end encrypted — the platform has to read a message to answer it.

12. Children

Tephlo is a business tool and is not directed at children. Workspaces must not configure an assistant to collect personal data from children below the age of consent in their jurisdiction, and must not use the platform for services aimed at children without a lawful basis and appropriate safeguards. If you believe a child's data has reached the platform, tell us and we will work with the workspace to remove it.

13. Automated processing

Replies are generated automatically, and the platform decides automatically when to bring in a person — for example when it is not confident, when a topic is sensitive, when someone asks for a human, or when a customer is visibly frustrated. These decisions affect how a conversation is handled; they do not by themselves produce legal or similarly significant effects about a person, and workspaces are required by the [Acceptable Use Policy](#) not to use the assistant to make such decisions. Human handoff depends on the workspace having an available, confirmed team route; response times also depend on that team's staffing and operating hours.

14. Changes to this policy

The date at the top of this page is the date this text last changed. Material changes are notified to workspace administrators by email before they take effect. New sub-processors appear on the [sub-processor page](#) before they begin processing.

15. Contact and complaints

Write to hello@tephlo.com, or to the data protection contact at **[DATA PROTECTION CONTACT]**. Postal address: TEPHLO AI ENTERPRISES, 9A Kenneth Lekara Luka Street, Lokogoma, FCT, Nigeria.

If you are not satisfied with our answer, you can complain to the supervisory authority in your country. In Nigeria that authority is the Nigeria Data Protection Commission (NDPC). We would rather hear from you first, but you are not required to come to us before going to them.