

Tephlo legal documents

8 documents in one file: Terms of Service, Deleting your data, Privacy Policy, Data Processing Agreement, Sub-processors, Acceptable Use Policy, Cookies and local storage, Security overview.

Revision Draft 1.1 — factual content refresh

Exported 2026-10-04T09:56:17.530Z

Terms of Service — Version 3 · Effective 26 September 2026

Deleting your data — Effective [EFFECTIVE DATE]

Privacy Policy — Version 2 · Effective 6 September 2026

Data Processing Agreement — Version 2 · Effective 6 September 2026

Sub-processors — Effective [EFFECTIVE DATE]

Acceptable Use Policy — Version 2 · Effective 6 September 2026

Cookies and local storage — Effective [EFFECTIVE DATE]

Security overview — Effective [EFFECTIVE DATE]

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 2026-10-04T09:56:17.530Z. The pages are the authoritative text.

Contents

Terms of Service	5
1. The agreement	5
2. What the service is	5
3. Accounts and access	6
4. Data and roles	7
5. Acceptable use	7
6. Messaging channels and other third parties	7
7. What the AI assistant does and does not do	8
8. Availability, maintenance and support	8
9. Fees	8
10. Term, suspension and termination	9
11. What happens to your data when the agreement ends	9
12. Confidentiality	10
13. Intellectual property	10
14. Warranties and disclaimers	10
15. Liability	10
16. Indemnity	11
17. Changes	11
18. Governing law and disputes	11
19. General	11
Contact	11
Deleting your data	13
1. Which of these are you?	13
2. You messaged a business that uses Tephlo	13
3. You run a Tephlo workspace	14
4. What erasure actually removes	14
5. What is deliberately kept	14
6. How long it takes	14
7. Requests arriving through Meta	15
Contact	15
Privacy Policy	16
1. Two different roles, and why it matters	16
2. What is collected from people who message a business	16
3. What the assistant refuses to collect	18
4. What is collected from the people who run a workspace	18
5. Website visitors	18
6. Why we handle this data, and on what basis	19
7. How long data is kept	19
8. Who else receives this data	20
9. International transfers	21
10. Your rights, and how they work here	21
11. How the data is protected	22
12. Children	22
13. Automated processing	22

Contents (continued)

14. Changes to this policy	22
15. Contact and complaints	22
Data Processing Agreement	23
1. Scope	23
2. Roles of the parties	23
3. Processing on your instructions	24
4. Our obligations	24
5. Your obligations	24
6. Security measures	25
7. Sub-processors	25
8. Data subject requests	25
9. Personal data breaches	26
10. Audit and information	26
11. International transfers	26
12. Deletion and return of data	27
13. General	27
Annex I — Details of the processing	27
Annex II — Technical and organisational measures	28
Contact	30
Sub-processors	31
Core provider roles	31
Optional — off unless switched on	32
A note on messaging channels	33
What is deliberately absent from this list	34
Changes and objections	34
Contact	34
Acceptable Use Policy	35
Who this applies to	35
1. Illegal and harmful use	35
2. Deception and impersonation	35
3. Messaging conduct	36
4. Sensitive data	36
5. High-stakes and regulated decisions	36
6. Platform integrity	37
7. Content you upload	37
8. What the platform enforces itself	37
9. What happens if this policy is broken	38
10. Reporting abuse	38
Contact	38
Cookies and local storage	39
1. The public sites	39
2. The one preference cookie	39
3. The consoles	40

Contents (continued)

4. The embeddable chat widget	40
5. WhatsApp and Telegram	40
6. Controlling what is stored	41
7. Why there is no cookie banner	41
Contact	41
Security overview	42
Registrations on the public record	42
1. Keeping workspaces apart	43
2. Signing in	43
3. Encryption	43
4. Making sure a message really came from the channel	44
5. What ends up in logs	44
6. Audit trails	44
7. Files a customer sends	45
8. The AI layer	45
9. Deletion and retention	46
10. Backups and recovery	46
11. How changes reach production	46
12. What we do not claim	47
13. Reporting a security problem	47
Contact	48

Terms of Service

The agreement between your business and the company behind Tephlo. It covers what you get, what you are responsible for, and what happens when either side wants to stop.

Last updated 26 September 2026 · Revision Draft 1.1 — factual content refresh · Version 3 · Effective 26 September 2026

IN PLAIN LANGUAGE

You get a workspace that answers your customers on WhatsApp, Telegram and your website using knowledge you upload, and hands conversations to your team when a person is needed. You keep ownership of everything you put in.

- You are responsible for what your workspace says and does, for the accuracy of the knowledge you upload, and for your staff accounts.
- The assistant can be wrong. It uses workspace knowledge, business facts and permitted capabilities, can ask for clarification, and brings in a person when required. You must not use it where a wrong answer causes serious harm.
- We do not promise a specific uptime, and the service depends on WhatsApp, Telegram and an AI provider that can fail independently of us.
- You can export or delete a customer's data at any time, and deleting your workspace really deletes it — see section 11.
- Either side can end the agreement; we can suspend a workspace immediately if it is causing harm or breaking a channel provider's rules.

This summary is a reading aid. Where it differs from the text below, the text below is what applies.

1. The agreement

These terms are between TEPHLO AI VANGUARD SYSTEMS LTD, a private company limited by shares incorporated in the Federal Republic of Nigeria under the Companies and Allied Matters Act 2020 (registration number 9893569) with its registered office at 9A Kenneth Lekara Luka Street, Lokogoma, FCT, Nigeria (“we”, “us”), and the business that creates a Tephlo workspace (“you”). They take effect on 26 September 2026 or, if later, when you first create a workspace or accept them in the console.

They apply together with the [Acceptable Use Policy](#), the [Privacy Policy](#) and, where you handle personal data through the service, the [Data Processing Agreement](#). If you have signed a separate order form or written agreement with us, that document wins wherever it conflicts with these terms.

This is a business-to-business service. It is not offered to consumers, and you confirm you are entering this agreement in the course of a trade, business or profession.

2. What the service is

Tephlo is a hosted customer-support assistant. In summary, a workspace can:

- receive and reply to messages on WhatsApp Business, Telegram and an embeddable web chat widget on your own site, and an enabled commercial API integration;
- answer from a knowledge base you upload — documents and pasted text — and from your workspace profile and structured product catalogue, with source information where available;
- escalate a conversation to your team, with its history and the reason, and let an agent take over on the same channel;
- capture structured requests such as a callback, a complaint or a refund request for your team to action;
- save carts and submit potential-order requests for your team to review; a request is not a placed, paid or shipped order, and this workflow does not collect card numbers or arrange fulfilment;
- handle appointment requests and, where automatic confirmation is enabled, confirm a booking only after the selected time is reserved;
- give you a console with conversations, knowledge, analytics, business hours, escalation review and privacy tooling.

Some capabilities are switched off by default and only work when both the platform and your workspace enable them — document reading, optical character recognition on scanned PDFs, image understanding, voice-note transcription and spoken replies are examples. Provider configuration, workspace settings and other capability controls also apply. Customer memory is enabled by default in the software, can be disabled by the workspace and is subject to sensitivity controls. The [product documentation](#) describes what is available; nothing in it forms part of this agreement, and we may change or withdraw individual features (see section 17).

Where Smart Updates is enabled, customers can consent to updates on a topic. The platform records that consent, checks send eligibility and includes a STOP instruction. Contacting your business alone does not subscribe someone. AI-generated answers and extracted facts may still contain errors; review your published information and action settings.

We provide the service as software you access over the internet. We do not provide contact-centre staffing, and we do not act as your agent in any conversation.

3. Accounts and access

Your workspace is isolated from every other workspace: data is stored against your workspace identifier and every request is scoped to the account that made it. Within your workspace, you decide who has an account and what role they hold.

You are responsible for:

- keeping account credentials confidential, and for everything done through your accounts;
- removing access promptly when someone leaves your team — the console lets you disable or delete an account, and a workspace must always keep at least one active administrator;
- the channel credentials you connect (WhatsApp tokens, a Telegram bot token), which are stored encrypted and used only to operate your channels.

Multi-factor authentication is enforced for platform operator accounts in production and is available for your administrators. We strongly recommend enabling it: it is the single control that most reduces the damage of a stolen password.

You must not share one account between people, probe or test the security of the platform without our written permission, or attempt to reach another workspace's data.

4. Data and roles

You own the content you put into the service: your knowledge base, your catalogue, your configuration, and the conversations your customers have with your workspace. We claim no ownership of it.

For personal data inside those conversations you are the controller and we are the processor: you decide why it is collected and we handle it on your instructions. Those instructions are given through how you configure and use the service. The [Data Processing Agreement](#) sets out the detail and forms part of this agreement.

We use your content to provide and maintain the service, and to diagnose faults. We do not sell it, and we do not use your customers' conversations to train AI models. Aggregated, non-identifying operational statistics — message volumes, error rates, latency — are used to run and improve the platform.

Quality features are opt-in. Features that keep customer content for review or reuse — escalation review excerpts, approved answer suggestions, response caching, customer memory — are controlled per workspace and default to off, apart from customer memory. See the [Privacy Policy](#) for what each one stores.

5. Acceptable use

The [Acceptable Use Policy](#) forms part of these terms and lists what a workspace may not be used for. In short: nothing illegal, nothing designed to deceive, no unsolicited bulk messaging, no collecting card numbers or one-time codes through chat, no use as a substitute for emergency services or regulated professional advice, and nothing that breaks the rules of the messaging channel you are using.

You are responsible for your customers' experience of your workspace, including the instructions you give the assistant and the accuracy of the knowledge you upload. Incorrect material can lead to incorrect answers, and generated replies can still contain errors.

6. Messaging channels and other third parties

WhatsApp and Telegram are operated by third parties on their own terms. Where you connect your own WhatsApp Business account or Telegram bot, your relationship with that provider is yours, and their policies — including messaging limits, template rules and account bans — apply to you directly. We cannot restore access that a channel provider withdraws.

Two consequences are worth stating plainly. WhatsApp only allows free-form replies within 24 hours of a customer's last message; outside that window an approved template is required, so a delayed reply may not be deliverable. Channel readiness and health checks can block delivery or disable an affected connection, depending on the channel ownership and failure type.

The service also depends on the providers listed on the [sub-processor page](#), including an AI provider that generates answers. We choose these providers carefully and remain responsible to you for the

service, but we do not control their availability.

7. What the AI assistant does and does not do

The assistant answers from material you supply. It is built to say it is unsure and hand off to a person rather than invent an answer, and several limits are enforced in code rather than merely requested of the model: it will not ask a customer for a card number, PIN, one-time code or similar credential; it will not report an action as completed unless the underlying step actually completed; and phrases you add to a “never say” list are checked before a reply is sent, with the conversation escalated to a person instead.

Those are real safeguards, and they are not a guarantee of correctness. Language models make mistakes, retrieval can surface the wrong passage, and your own documents may be out of date. You must:

- review the knowledge you upload, and keep it current;
- test your workspace before pointing customers at it, and monitor the conversations it handles;
- not use the service to give medical, legal or financial advice that a qualified professional should give, to make decisions with legal or similarly significant effects about a person, or in any setting where a wrong answer risks physical harm.

Not an emergency service. The assistant must never be presented to customers as a route to emergency help. It replies only when the platform is running and a channel is working, and it is not monitored for emergencies.

8. Availability, maintenance and support

We aim to keep the service available and we monitor it, but we do **not** offer a service level agreement, an uptime percentage, or service credits. Anyone telling you otherwise is mistaken.

We may take the platform into maintenance mode for planned work, which shows a notice to signed-in users and can pause AI replies. We will give reasonable notice of planned maintenance where we can, and we may act without notice where there is a security or stability risk.

Support is provided by email at hello@tephlo.com during business hours in West Africa Time (WAT). Response targets, if any, are those in your order form; there are none by default.

9. Fees

Plans and prices are published on our pricing page and may be set out in an order form. Available billing behavior depends on the deployment:

- **Workspace plan payments.** Where Paystack billing is enabled, the console can open a hosted checkout for a workspace plan. Payment details are entered on the provider's page. The platform stores transaction records and, where supplied for renewal, a reusable authorization reference and limited payment-method display details; it does not receive full card numbers through that checkout flow. Payment and renewal terms remain those presented to you and agreed in your order form.

- **Usage and limits.** The platform meters usage and can enforce plan allowances when the relevant platform and entitlement controls are enabled. Hard caps, grace units and any permitted overage depend on the active plan and configuration. A usage counter alone is not proof that a payment was taken.

Paying for your workspace plan is separate from a customer submitting a saved cart to your business. A potential-order request does not take payment for that cart or confirm an order.

Fees are exclusive of taxes, which you pay in addition where they apply. Late payment terms, currency and payment method are as stated in your order form; unless it says otherwise, all payments are due upfront in Nigerian Naira (NGN), by recurring debit-card authorization, bank transfer or an electronic payment gateway. We may change prices for a renewal term on 30 days' notice.

10. Term, suspension and termination

The agreement runs from when you accept it until it is ended. Unless your order form says otherwise, either side may end it for convenience on 30 days' written notice, and fees already paid for a period are not refunded for the unused part of it.

Either side may end the agreement immediately if the other:

- materially breaches it and does not fix the breach within 30 days of being asked to; or
- becomes insolvent or stops trading.

We may suspend a workspace immediately — including disabling its channels — where we reasonably believe it is being used in breach of the [Acceptable Use Policy](#), is putting the platform or other customers at risk, is breaching a channel provider's rules, or where a channel provider requires it. We will tell you why, and lift the suspension when the cause is resolved. Suspension is also a technical prerequisite for deleting a workspace, so that a deletion cannot race live traffic.

11. What happens to your data when the agreement ends

Before the end, you can export your data through the console and the API, including a full export of an individual customer's conversations.

On termination we delete your workspace. That deletion is real and irreversible: it removes conversations and message history, escalations, captured tickets, knowledge sources and their indexed chunks, product catalogue, customer memory, channel configuration and encrypted credentials, staff accounts, usage records, uploaded media and their extracted text, and the queued work belonging to the workspace. Files in object storage are removed by a job that commits with the deletion, so a crash cannot leave them orphaned.

Two things deliberately survive, and you should know about both:

- **Platform audit records.** A small number of audit tables record that an administrative action happened — a deletion, a channel credential change, a routing change — with who did it and why. They hold no message content, and two of them cannot be modified even by us, because the database itself refuses. They are the evidence that the platform was operated correctly.
- **Backups.** Encrypted database backups are kept on a rolling schedule and are not individually edited. Deleted data disappears from backups as those backups age out and are pruned, not at the moment of deletion.

Unless you ask us to delete sooner, we will delete a terminated workspace within 30 days of the end of the agreement, so that a mistaken termination can be reversed.

12. Confidentiality

Each side may receive information from the other that is marked confidential or that a reasonable person would treat as confidential. Each will use it only to perform this agreement, protect it with at least reasonable care, and disclose it only to people who need it and are under similar obligations. This does not apply to information that is public through no fault of the receiver, was already known, is independently developed, or must be disclosed by law — in which case the receiver gives notice where it lawfully can.

13. Intellectual property

We own the platform, its software, and everything we supply with it. You get a non-exclusive, non-transferable right to use it during the term for your own business. You own your content, and you grant us the licence needed to host, process and display it in order to run the service for you.

You may not copy, reverse engineer, resell or offer the platform as your own service, or use it to build a competing product, unless a separate written agreement says you may. Feedback you give us can be used freely.

14. Warranties and disclaimers

Each side warrants that it has the authority to enter this agreement. We warrant that we will provide the service with reasonable skill and care.

Beyond that, and to the extent the law allows, the service is provided “as is”. We do not warrant that it will be uninterrupted or error-free, that AI-generated answers will be accurate or complete, or that it will meet any particular regulatory requirement that applies to your business. Nothing in this section limits rights that cannot be limited under the law of the Federal Republic of Nigeria.

You warrant that you have the right to upload the content you upload and to have the personal data in it processed as described in the [DPA](#).

15. Liability

Neither side excludes liability for death or personal injury caused by negligence, for fraud, or for anything else that cannot lawfully be excluded.

Subject to that, neither side is liable for loss of profit, revenue, goodwill, anticipated savings, or for indirect or consequential loss. Each side's total liability arising out of this agreement is capped at the total amounts actually paid by you to us in the six (6) months immediately preceding the event giving rise to the liability.

Given what the service does, one allocation is stated explicitly: you are responsible for the content of your knowledge base and for reviewing what your workspace tells your customers, and we are not liable for the consequences of an answer that faithfully reflects inaccurate material you supplied.

16. Indemnity

You will defend and indemnify us against third-party claims arising from your use of the service in breach of this agreement or the [Acceptable Use Policy](#), or from content you uploaded. We will defend and indemnify you against third-party claims that the platform itself infringes their intellectual property rights, provided you tell us promptly and let us run the defence.

17. Changes

We may change the service, and we may change these terms. For material changes to these terms we will give workspace administrators 30 days' notice by email before they take effect; if you do not accept a material change, you may end the agreement before it applies. Changes required by law or to address a security risk may take effect immediately.

We may add, change or withdraw individual features. We will not remove a feature you rely on without notice unless a provider or a security issue forces it.

18. Governing law and disputes

This agreement is governed by the laws of the Federal Republic of Nigeria. Any dispute, controversy or claim arising out of or relating to it is referred to and finally resolved by binding private arbitration in Abuja, Federal Capital Territory, administered by a single arbitrator in accordance with the Arbitration and Mediation Act of Nigeria. The costs of the arbitration, including the arbitrator's fees and the venue, are borne as the arbitrator determines or otherwise shared equally by the parties, and each side bears its own legal representation fees. Before starting arbitration, each side agrees to raise the issue in writing and to try in good faith to resolve it for 30 days.

19. General

- **Assignment.** Neither side may assign this agreement without the other's consent, except to a successor of its business.
- **Notices.** Notices to you go to the email addresses of your workspace administrators; notices to us go to hello@tephlo.com and to the registered address above.
- **Force majeure.** Neither side is liable for failure caused by events outside its reasonable control, including provider outages, network failures and government action.
- **Severability.** If a provision is unenforceable, the rest continues in force.
- **No waiver.** Not enforcing a right once does not waive it.
- **No partnership.** Nothing here creates a partnership, agency or employment relationship.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at hello@tephlo.com. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: TEPHLO AI VANGUARD SYSTEMS LTD, 9A Kenneth Lekara Luka Street,
Lokogoma, FCT, Nigeria.

Deleting your data

Two different people read this page, and they need two different answers. Start with which one you are.

Last updated 6 September 2026 · Revision Draft 1.1 — factual content refresh · Effective [EFFECTIVE DATE]

IN PLAIN LANGUAGE

If you messaged a business that uses Tephlo, that business decides what happens to your data — we hold it on their instructions. Ask them, and they can erase you from their dashboard straight away. If you cannot reach them, write to us at hello@tephlo.com and we will pass the request on and follow up.

If you run a Tephlo workspace, you can erase any individual yourself from Privacy Center, and you can close the whole workspace, which removes everything in it.

- Erasure is scoped to a workspace, channel and exact identity — a phone number, Telegram id, web id or API customer id — and it is irreversible.
- It removes the conversation history, the stored profile and the live session state, not just a flag on a row.
- A record that a deletion happened is kept. It contains no message content and no readable identifier.

1. Which of these are you?

Tephlo is used by businesses to answer their own customers. That makes two groups of people whose data we hold for entirely different reasons, and the route to erasure is different for each. The [Privacy Policy](#) explains the split in full; the short version is that for customer conversations the business is the controller and we are its processor, and for workspace accounts we are the controller.

2. You messaged a business that uses Tephlo

Ask that business. They can erase you immediately from their dashboard, without contacting us, and they do not need our permission — that is how it is meant to work, because the conversation is theirs.

If you do not know who they are, cannot reach them, or would rather not contact them, write to hello@tephlo.com with the phone number, Telegram user id, web session or API customer id you used, and which business you were messaging. We will identify the workspace, pass the request to them, and tell you what happened. We will not erase a business's records on our own initiative, because we are not the ones who decided to collect them — but we will not let a request go unanswered either.

You do not need an account with us to make this request, and there is no charge for it. The [privacy request form](#) sends it to the team as a tracked request.

3. You run a Tephlo workspace

Two different things, depending on what you want gone.

One person's data. Open **Privacy** in your dashboard, find the customer, and confirm the deletion by re-typing their exact identifier. The re-typing is deliberate: erasure cannot be undone, and a mistyped identifier would otherwise silently erase the wrong person or nobody at all. The same screen exports their data if you need to give them a copy first.

The entire workspace. Ask us to close it. That removes the workspace and everything held for it — conversations, customers, knowledge, catalogue, media and staff accounts. Your channel connections are released, so your WhatsApp Business account and telephone number remain yours and can be connected elsewhere.

4. What erasure actually removes

Deleting one person is not a flag on a row. For that identity we remove the conversation history and message content, the customer profile and anything remembered about them, support tickets and structured care cases, customer-scoped cached answers and review excerpts, and the live session state held outside the database. It also removes Smart Updates consent and send history, saved carts and their lines, potential-order requests, appointments, and related delivery records and unsent drafts. Export and erasure have different inventories: operational records can be erased without being separately included in the customer export.

Erasure also takes a lock on that identity while it runs. Anything already in flight either lands before the deletion and is removed with everything else, or arrives afterwards and is refused. There is no window in which a message slips in behind an erasure and quietly recreates the person.

5. What is deliberately kept

A record that a deletion happened: when, which workspace, on whose authority, and the stated reason. It holds no message content and no readable identifier — the subject is stored as a keyed digest, so the record proves the erasure occurred without preserving the person it was about.

That record exists because an erasure you cannot demonstrate is not much use to anybody: it is what lets a business show a regulator, or you, that your request was carried out. Deleting it would remove the only evidence that your data was deleted.

Backups are another exception. A completed erasure removes the covered live records. A backup made before that request can retain a copy until it is pruned under the deployment's backup schedule. The operator must reapply erasures when restoring a backup; the live erasure action does not edit old backup files.

6. How long it takes

A successful dashboard erasure removes the covered database records as part of that operation. External storage cleanup can require a worker; its completion depends on that worker running. A request that comes to us and has to be routed to a business depends on them; we acknowledge within a few working days and chase if they do not act. Where the law that applies to you sets a

deadline, that deadline governs and we work to it.

7. Requests arriving through Meta

If you reached this page from Facebook or WhatsApp settings, you are probably trying to remove data associated with a Meta login. Tephlo uses Meta sign-in only so a business can connect its own WhatsApp Business account to us; we do not build profiles from it and we do not receive your Facebook friends, posts or activity.

Disconnecting Tephlo from a WhatsApp Business account revokes our access to it. To have the associated records erased as well, use the route in section 2 or 3 above, or write to us — either reaches the same people.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.

TEPHLO AI ENTERPRISES is the operator of Tephlo. If you are unhappy with how a request was handled you can complain to the Nigeria Data Protection Commission (NDPC), and if your own country has a different authority you may complain to that one instead. Written requests can also be sent to our registered address, which is on the [Privacy Policy](#).

Effective from **[EFFECTIVE DATE]**.

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

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.

Data Processing Agreement

For businesses using Tephlo to handle their customers' personal data. You decide what is processed and why; we process it on your instructions, and this document says exactly how.

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

IN PLAIN LANGUAGE

You are the controller of your customers' personal data. We are your processor. This document is the contract for that relationship.

- We process your customers' data only to run the service for you, and not for our own purposes. We do not use it to train AI models.
- Annex II lists the security measures that are actually implemented. Where a measure is organisational rather than technical, and we cannot verify it from the code, it is marked as a gap rather than claimed.
- The [sub-processor list](#) is part of this agreement. You are notified before a new one starts, and you can object.
- We tell you about a personal data breach without undue delay after becoming aware of it, and help you meet your own notification duties.
- On request or on termination we delete your data. Section 12 says exactly what that removes, what it does not, and what happens with backups.

1. Scope

This Data Processing Agreement ("DPA") forms part of the [Terms of Service](#) between TEPHLO AI ENTERPRISES ("Processor", "we") and the business operating a Tephlo workspace ("Controller", "you"). It applies to personal data we process on your behalf when you use the service, and takes effect on **[EFFECTIVE DATE]** or when you first accept the Terms, whichever is later.

Terms such as "personal data", "processing", "controller", "processor", "data subject" and "personal data breach" carry the meanings given to them in the Nigeria Data Protection Act 2023. Where that law uses different labels, the equivalent concepts apply.

2. Roles of the parties

You are the controller of the personal data in your workspace: your customers' messages and identifiers, the knowledge you upload, and the configuration you set. We are the processor. Where you are yourself acting as a processor for someone else — an agency running a workspace for a client — we are a sub-processor, and your own controller's instructions reach us through you.

We are an independent controller for a narrow, separate set of data: the accounts of the people who administer your workspace, our billing and support records, and platform security and audit records. That processing is described in the [Privacy Policy](#), not here.

3. Processing on your instructions

We process personal data only on your documented instructions, including for international transfers, unless required otherwise by law — in which case we tell you before processing, unless that law forbids it.

Your instructions are:

- this DPA and the Terms of Service;
- your configuration of the service — the channels you connect, the knowledge you upload, the retention and privacy settings you choose, and the optional features you enable;
- your use of the console and the API, including exports and deletions you request;
- any further written instructions we agree with you.

We will tell you if, in our opinion, an instruction infringes applicable data protection law. We are not obliged to give you legal advice, and we may decline to act on an instruction that would put other customers or the platform at risk.

4. Our obligations

- **Purpose limitation.** We use your customers' personal data only to provide, secure and maintain the service for you. We do not sell it, do not use it for advertising, and do not use it to train AI models.
- **Confidentiality.** Access is limited to personnel who need it to operate the platform, under an obligation of confidentiality.
- **Security.** We implement the technical and organisational measures in Annex II, and will not materially weaken them during the term.
- **Assistance.** We assist you, taking into account the nature of the processing and the information available to us, with data subject requests (section 8), with breach notification (section 9), and with data protection impact assessments and prior consultations, for which this document and the [Security Overview](#) are the primary input.
- **Records.** We maintain records of the categories of processing carried out on your behalf, and make them available on request.

5. Your obligations

- You are responsible for having a lawful basis for the processing you instruct, and for telling your customers what happens to their data — including that an automated assistant answers them and that a person may take over.
- You configure retention, privacy settings and optional features. The defaults are documented; what applies to your workspace is what you set.
- You do not upload special categories of personal data (health, biometric, political, religious and similar) or children's data without a lawful basis and appropriate safeguards, and you keep your knowledge base free of personal data that does not need to be there.

- You keep your accounts secure and remove access promptly when someone leaves your team.

6. Security measures

Annex II lists the measures in force. They are described at a level a reviewer can check, and where a measure is organisational — personnel vetting, security training, formal policy review — it is marked as an open item rather than asserted, because this document should not be the first place such a claim appears.

No certification. We hold no SOC 2 report and no ISO 27001 certificate, and no independent auditor has assessed these measures. If your procurement process requires one, tell us before you sign rather than after.

7. Sub-processors

You give general authorisation for us to engage sub-processors. The current list, with the purpose, data categories and location of each, is published at [legal / sub-processors](#) and forms part of this DPA.

Before a new sub-processor starts processing your data, we will update that page and notify workspace administrators by email, giving **[SUB-PROCESSOR NOTICE PERIOD]** notice. If you reasonably object on data protection grounds within that period, we will work with you on an alternative; if none is available, you may terminate the affected part of the service without penalty for the remainder of the term.

We impose data protection obligations on each sub-processor that are no less protective than those in this DPA, and remain responsible to you for their performance.

One nuance about messaging providers. Where you connect your own WhatsApp Business account or Telegram bot, that provider is processing under your own relationship with them, not as our sub-processor. Where you use a number the platform shares, our relationship with the provider applies. Both cases are marked on the sub-processor page.

8. Data subject requests

The service gives your administrators the tools to answer most requests without involving us:

- **Access and portability.** An export produces a structured record of one customer's data with your workspace — conversations, the full transcript, escalations, sessions on shared numbers, routing records, memberships, remembered facts, tickets and care cases, customer-scoped cached answers, Smart Updates consent and send history, carts, potential-order requests and appointments. Review excerpts are represented by metadata; their message text is already disclosed in the transcript.
- **Erasure.** A deletion permanently removes that customer's conversations and messages, escalations, transcripts of media in those conversations, sessions, routing records, memberships, queued outbound messages and in-flight processing records, and clears their short-lived working state. Anything already accepted for processing before the deletion is suppressed rather than delivered afterwards.

- **Rectification and restriction.** Handled by editing or removing the underlying records in the console.

Every export and deletion is recorded in an audit trail that identifies the subject by a salted hash rather than by their identifier, so the audit itself does not recreate the data you just erased.

Scope of the per-customer erasure. It removes the conversations and messages, the escalations raised from them, the routing, session and membership records, queued deliveries, the entries stored by the customer-memory feature, and the tickets captured for your team with their summaries, care cases, customer-scoped cached answers, review excerpts, Smart Updates consent and send records, carts, potential-order requests and appointments. Related operational rows and withheld reply drafts are erased but are not all separately disclosed in an export. The deletion audit itself is retained, identifying the subject only by a salted hash, so that the erasure can be evidenced without recreating what it removed.

If a data subject contacts us directly about your workspace, we will not respond substantively — we will refer them to you and tell you promptly.

9. Personal data breaches

We notify you without undue delay after becoming aware of a personal data breach affecting personal data we process for you. The notice will describe, as far as we know at the time: the nature of the breach and the categories and approximate number of data subjects and records affected, the likely consequences, the measures taken or proposed, and a contact point. Where the full picture is not available immediately, we send what we have and follow up as we learn more, rather than delaying the first notice.

We will assist you with your own notification obligations to a supervisory authority or to affected individuals. Notifying us of a breach in your own systems, or a compromised staff account, is your responsibility — write to hello@tephlo.com.

10. Audit and information

On request, and no more than once a year unless a supervisory authority or a breach requires otherwise, we will provide the information reasonably necessary to demonstrate compliance with this DPA: this document, the [Security Overview](#), our answers to a reasonable security questionnaire, and a description of any material change to the measures in Annex II.

Because there is no third-party certification to hand you instead, an on-site or third-party audit is available on the terms in **[AUDIT TERMS]** — scope, notice, confidentiality and costs to be agreed. We will not give access to another customer's data or to anything that would compromise the security of the platform.

11. International transfers

The service is delivered using providers that operate internationally, and the AI provider may route a request to a model provider in another country. Personal data will therefore be transferred across borders.

The transfer mechanism relied on for each sub-processor is **[TRANSFER MECHANISM]**, to be confirmed by counsel before this document is used as a signed agreement. We will not claim a mechanism is in place before it is.

12. Deletion and return of data

You can export your data at any time during the term through the console and the API. On termination, or on your written request, we delete the personal data we process for you.

Deleting a workspace removes:

- conversations, messages, escalations and captured tickets;
- knowledge sources and their indexed chunks, the product catalogue, remembered customer facts, care cases, Smart Updates records, carts, potential-order requests and appointments;
- channel configuration and encrypted credentials, staff accounts and their credential records, invites and pending signups;
- uploaded media and extracted text, including the files themselves in object storage, removed by a cleanup job that commits together with the deletion; physical removal depends on that worker and the storage provider succeeding;
- usage records, analytics rows, email delivery records for the workspace, and queued work belonging to it.

Two categories survive, deliberately:

- **Platform audit records** showing that administrative actions took place — a deletion, a credential change, a routing change — with the actor and the stated reason. They contain no message content, and two of these tables are protected against modification by the database itself.
- **Encrypted backups**, which are not edited individually. Deleted data leaves the backup set as backups are pruned on their schedule. Until then it remains encrypted and is restored only in a disaster, after which the deletion is reapplied.

Unless you instruct otherwise, a terminated workspace is deleted within **[POST - TERMINATION RETENTION PERIOD]**, so a mistaken termination can be undone. We will confirm the deletion in writing on request.

13. General

This DPA runs for as long as we process personal data for you. If it conflicts with the Terms of Service on the treatment of personal data, this DPA prevails. Liability under this DPA is subject to the limitations in the Terms. If a provision is held invalid, the rest remains in force, and the parties will agree a valid replacement with equivalent effect.

Annex I — Details of the processing

Subject matter

Provision of an AI customer-support assistant and agent console across WhatsApp, Telegram, web chat and enabled commercial API integrations.

Duration

The term of the Terms of Service, plus the deletion window in section 12.

Nature and purpose

Receiving and storing customer messages; retrieving relevant passages from your knowledge base; generating replies through an AI provider; routing conversations to your team; capturing requests such as callbacks and complaints; managing saved carts, potential-order requests and appointments; producing analytics for you; and, where permitted by the relevant controls, interpreting images, transcribing voice notes, reading documents, generating spoken replies and remembering stated customer preferences. Smart Updates records topic consent and preferences and sends eligible approved updates.

Categories of data subjects

The customers and prospective customers who message your workspace, and any individuals they mention in a message.

Categories of personal data

Channel identifiers (a phone number on WhatsApp, a user id on Telegram, a randomly generated visitor id in web chat, or the identifier supplied by an API integration); the content of messages sent and received, which may contain anything a customer chooses to write; conversation and escalation metadata; captured request summaries and care-case records; stated preferences where customer memory is allowed; Smart Updates consent and send history; saved-cart selections and price snapshots, potential-order requests and appointment details; transcripts, images, extracted document text and spoken reply content where the relevant features are enabled.

Special categories

Not intended. The service is not designed for them and you agree not to instruct their processing without a lawful basis and safeguards. A customer may nonetheless volunteer such information in a free-text message; the platform treats it as ordinary message content.

Frequency

Continuous, for as long as your channels are live.

Annex II — Technical and organisational measures

These are implemented in the platform today. The [Security Overview](#) explains each in plain language.

1. **Workspace isolation.** Workspace records carry their ownership, and access is checked against the authenticated principal or authorised channel/API configuration. Supplying a record identifier does not itself grant access to it. Conversations are keyed by workspace, channel, receiving identity and customer.
2. **Access control.** Role-based permissions; authorisation re-checked against the database on every request, so a disabled account, a changed role or a suspended workspace takes effect immediately; short-lived access tokens; refresh tokens that rotate and whose whole family is revoked if an old one is replayed; account lockout after repeated failures.
3. **Multi-factor authentication.** Enforced for platform operator accounts in production and available to your administrators; the authenticator secret is stored encrypted.

4. **Encryption in transit.** HTTPS throughout; production configuration refuses to start with a non-HTTPS provider endpoint, requires TLS to the cache, and requires an encrypted connection to a managed database.
5. **Encryption at rest.** Channel credentials, authenticator secrets, invite payloads, email bodies, staged uploads, media job payloads and extracted document text are encrypted with authenticated symmetric encryption; supported key rings allow staged key rotation. Passwords are stored only as hashes, not in readable form.
6. **Message authenticity.** Inbound WhatsApp webhooks are verified by HMAC signature over the original request body. Activation and health controls depend on the channel and its ownership; a single send failure does not automatically disable a platform-shared channel.
7. **Log hygiene.** Email addresses, phone numbers, URLs, tokens, signatures and stack traces are redacted at the logging boundary. Audit and analytics records store categories and counts, never message content.
8. **Audit trails.** Administrative actions — credential changes, routing changes, privacy exports and deletions, platform configuration — are recorded with the actor, target and stated reason. Two audit tables are append-only at the database level: an update or delete is rejected by the database, not merely by the application.
9. **File safety.** Documents a customer sends are scanned for malware before anything opens the bytes, and the scan fails closed — no scanner means no parsing. Executables, archives, macro-enabled documents, encrypted files and PDFs containing active content are rejected outright; extraction runs under hard bounds; infected files are deleted unparsed and can never be downloaded through the application.
10. **Private storage.** Uploaded media is stored under opaque, backend-generated keys partitioned per workspace, with no public access and no pre-signed links; customer filenames never appear in storage keys.
11. **Sensitive-data minimisation.** The assistant refuses to request payment instruments, credentials, bank details and government identity numbers, in English, French and Spanish. Text kept for optional quality review is masked for emails, phone numbers, long numbers and labelled codes before storage.
12. **Resilience.** Work is staged in a durable outbox with retries, backoff and a dead-letter queue; duplicate deliveries are collapsed; failing providers are isolated by circuit breakers so one outage does not cascade.
13. **Backups and recovery.** An encrypted backup procedure, restore procedure and integrity checks are documented. Their operation, schedule and key management must be confirmed for the deployment.
14. **Data lifecycle.** Time-limited caches and working state, per-customer erasure and workspace deletion. Current memory extraction sets a 180-day expiry; expired facts are excluded on read. Physical removal by retention sweeps or object-storage cleanup depends on the relevant worker being enabled and running.
15. **Change management.** Version-controlled code and database migrations, automated checks on every change (linting, type checking, tests, migration checks), and staged rollout of risky features behind flags that default to off.
16. **Organisational measures.** Personnel confidentiality undertakings, background checks, security training, formal policy review and an incident response plan owned by a named person:

[ORGANISATIONAL MEASURES TO BE CONFIRMED]. These are not claimed here until the company can evidence them.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.

Sub-processors

The third parties involved in running Tephlo, what each one is used for, which data reaches it, and where it operates. This list is part of the Data Processing Agreement.

Last updated 6 September 2026 · Revision Draft 1.1 — factual content refresh · Effective [EFFECTIVE DATE]

IN PLAIN LANGUAGE

Running an AI assistant on WhatsApp and Telegram means other companies are involved: one generates the answers, two carry the messages, one hosts the database, one sends email, one stores files. The roles and supported provider configurations are listed below. The actual providers, regions and agreements must be checked against the deployment before this draft is treated as a complete register.

- Several are **off by default** — voice transcription, spoken replies, image reading, web search and error monitoring only apply if they have been switched on.
- No analytics, advertising or tracking company appears on this list, because none is used.
- These providers operate internationally. We cannot promise data stays in one country, and this page says where each one is based rather than implying more precision than exists.

Core provider roles

These describe supported and documented configurations, not proof that each provider is active for every workspace. Endpoint settings and the operator's provider agreements determine the actual recipients.

IN USE — ALL WORKSPACES

Provider	Purpose	Data categories	Location
OpenRouter	Generating the assistant's replies. Every AI answer goes through this provider, which routes the request on to the selected model provider.	The customer's message and recent conversation, the passages retrieved from the workspace's knowledge base, and the workspace's instructions to the assistant. Whatever a customer writes in a message can therefore reach this provider.	United States, and the country of whichever model provider it routes to.
OpenAI	Turning knowledge base text and customer questions into search vectors, so an answer can be found in the right document.	Knowledge base text uploaded by the workspace, and the text of customer questions used to search it. This is the default. A deployment can be configured to take search vectors from OpenRouter instead, in which case this row does not apply.	United States.

Meta Platforms (WhatsApp Business Cloud API)	Receiving and delivering WhatsApp messages. Required for the WhatsApp channel.	Customer phone numbers and the full content of messages in both directions.	Global infrastructure operated by Meta. See the note on channels below — with your own WhatsApp number, Meta is your provider, not our sub-processor.
Telegram	Receiving and delivering Telegram messages. Required for the Telegram channel.	Telegram user identifiers and the full content of messages in both directions.	Global infrastructure operated by Telegram.
Cloudflare	Object storage for files customers send and for brand assets; hosting for the public marketing, documentation and legal sites.	Uploaded files and their extracted text; brand images. The public sites involve only ordinary connection data.	Storage region set when the bucket was created: Cloudflare R2, automatic placement. Site hosting is distributed.
Heroku (Salesforce)	Application hosting, the managed PostgreSQL database, and the managed key-value store used for queues and short-lived state.	All application data at rest, including conversations, knowledge bases and accounts.	Region chosen at provisioning: Europe (Ireland). Confirm the live deployment before this page is published as final; the platform can also be self-hosted, in which case this row is replaced by the operator's own infrastructure.
Mailtrap	Delivering operational email: signup verification, team invitations, password resets and escalation alerts.	Staff names and email addresses, and the content of those emails. Escalation alerts carry no customer message content — they say a conversation needs a person and link to the console.	Endpoint configured per deployment (EU or US): [EMAIL REGION]. A workspace deployment may instead use its own SMTP provider.

Optional — off unless switched on

These are disabled by default. Each requires the platform to be configured for it and, where it affects conversations, the workspace to opt in. If they are off for your workspace, no data reaches them.

OPTIONAL — ENABLED PER DEPLOYMENT OR PER WORKSPACE

Provider	Purpose	Data categories	Location
OpenAI (speech to text)	Transcribing voice notes so the assistant can answer them.	The audio of a voice note, and the resulting transcript.	United States.
Paystack (workspace plan billing)	Hosted checkout, payment verification and supported renewals for workspace plans, where billing is enabled. This does not process payment for a customer's saved cart.	Payer email, amount, currency, payment references and transaction metadata. Payment details are entered directly with the provider; the platform can retain a reusable authorization reference and limited payment-method display details returned by it.	[BILLING PROVIDER REGION AND CONTRACTUAL ROLE]. Confirm against the configured account and agreement.
Configured speech provider (spoken replies)	Turning an approved assistant reply into audio, where spoken replies and a compatible speech endpoint are enabled.	The approved reply text, selected voice and language-related settings, and the resulting audio. This can contain customer information already present in the reply.	[SPEECH PROVIDER AND REGION]. The endpoint is configurable; an OpenAI-compatible interface does not establish which company receives the request.
Configured image provider (default: OpenAI)	Interpreting customer images and, when separately enabled, reading scanned PDF pages without selectable text.	Customer images or rendered pages of a customer document.	[IMAGE PROVIDER AND REGION]. Confirm the configured endpoint and provider before final publication.
Tavily	Web search, for question types a workspace has allowed the assistant to look up.	A search query derived from the customer's question.	United States.
Sentry	Error monitoring — collecting crash reports so faults get fixed.	Error type, stack trace and technical context. Message content, identifiers, tokens and addresses are redacted before logging.	Region of the configured project: [ERROR MONITORING REGION].
Telegram Gateway	Delivering password-reset codes to a staff member's phone, where they have chosen that method.	A staff phone number and a short one-time code.	Global infrastructure operated by Telegram.

A note on messaging channels

Who contracts with WhatsApp and Telegram matters. Most workspaces connect their own WhatsApp Business account or their own Telegram bot. In that case the provider processes under your relationship with them and their terms apply to you directly — they are not our sub-processor, and we cannot influence what they do with the message. Where a workspace uses a

number the platform shares between businesses, the relationship is ours and the row above applies. Either way, the provider necessarily sees every message sent over its own channel.

What is deliberately absent from this list

- **No analytics provider.** The marketing, documentation and legal sites load no analytics script — their content security policy blocks third-party scripts entirely.
- **No advertising or tracking network,** no advertising pixels, and no data brokers.
- **No third-party font service.** Public application scripts and fonts are served by the site. Published logos and icons may load from explicitly configured API or asset-hosting origins.
- **No cart-payment processing.** The built-in potential-order flow records a customer request for business review. The optional plan billing provider above is used to pay for the workspace itself.
- **No AI model training on your data.** Conversations are sent to a provider to generate a reply; they are not contributed to model training by us.

Changes and objections

Before a new sub-processor begins processing customer data, this page is updated and workspace administrators are notified by email, with **[SUB-PROCESSOR NOTICE PERIOD]** notice. If you object on reasonable data protection grounds, section 7 of the [Data Processing Agreement](#) sets out what happens next.

Verify before you rely on this. Provider locations and configuration change. The rows above describe how the platform is built and which providers it is designed to use; the regions marked as gaps must be confirmed against the live deployment before this page is treated as a final compliance record.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.

Acceptable Use Policy

An assistant that speaks to the public on your behalf can do real damage if it is pointed at the wrong job. This is the list of jobs it must not be pointed at.

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

IN PLAIN LANGUAGE

Use Tephlo to answer your customers honestly from your own information. Do not use it to deceive people, to spam them, to collect their card details, or to stand in for a professional or an emergency service.

- Some of these limits are enforced by the software, not just by this policy — section 8 lists which.
- Breaking the rules of WhatsApp or Telegram is also breaking this policy, and those providers can ban your account regardless of what we do.
- Serious or repeated breaches lead to suspension of the workspace and, if unresolved, termination.

Who this applies to

This policy applies to every workspace, every account in it, and everyone acting through it — including the assistant, since what it says is said on your behalf. It forms part of the [Terms of Service](#). Where a channel provider imposes stricter rules, theirs apply as well.

1. Illegal and harmful use

Do not use a workspace to:

- break the law of any country where you operate or where your customers are;
- sell or promote goods and services you are not permitted to sell, or that require a licence you do not hold;
- harass, threaten or defame anyone, or produce content that sexualises children, incites violence, or promotes self-harm;
- distribute malware, phishing links, or content designed to compromise someone's device or accounts;
- facilitate fraud, money laundering, or the evasion of sanctions or export controls.

2. Deception and impersonation

- Do not configure the assistant to claim it is a human being when someone sincerely asks. It may have a name and a personality; it must not deny being an automated assistant.
- Do not impersonate another business, a public body, or a person, or use a brand you have no right to use.

- Do not instruct the assistant to state things you know to be false — about prices, availability, guarantees, legal rights, or what your business will do next.
- Do not use it to generate fake reviews, testimonials or endorsements.

3. Messaging conduct

- Do not send unsolicited bulk or marketing messages. Smart Updates, where enabled, sends approved topic updates to customers with recorded consent. An incoming conversation alone does not subscribe someone, and the feature is not permission to contact an arbitrary list of people.
- Honour opt-outs immediately, and keep to the messaging rules, template requirements and quality ratings of the channel you use.
- Do not use the assistant to continue contacting someone who has asked you to stop.
- Do not connect a channel or number you are not authorised to operate.

4. Sensitive data

- Do not use the assistant to collect card numbers, CVV codes, PINs, passwords, one-time passcodes, full bank account details or government identity numbers. Chat is not a payment terminal and not a credential form.
- Do not configure it to request special categories of personal data — health, biometric, racial or ethnic origin, political opinions, religious beliefs, sex life or sexual orientation — without a lawful basis and appropriate safeguards.
- Do not direct it at children below the age of consent in your jurisdiction, or use it in services aimed at them, without a lawful basis and safeguards.
- Do not upload knowledge documents containing personal data that has no reason to be in an answerable knowledge base — staff lists, customer databases, internal notes about named individuals.

5. High-stakes and regulated decisions

The assistant uses your knowledge, business facts and enabled capabilities. That is useful for “what are your opening hours” and dangerous for “should I take this medication”. Do not use it:

- as an emergency, crisis or safety-of-life service, or in any way that suggests it can summon help;
- to give medical diagnosis or treatment advice, legal advice, or personalised financial or investment advice that a qualified professional should give — explaining your own published policies, products and processes is fine;
- to make or automate decisions that produce legal or similarly significant effects for a person: credit, insurance, employment, housing, benefits, or access to essential services;
- to confirm a payment, verify an identity, or authorise a transaction on the strength of a chat message or an uploaded document.

If a conversation is heading somewhere serious, route it to a person. The console lets you list situations that should always go to a human. Other uncertainty may lead to clarification, an allowed next step or escalation according to the workspace controls. These checks do not eliminate AI errors; review your configuration and keep a person available for cases that require judgment.

6. Platform integrity

- Do not attempt to access another workspace's data, or to enumerate customers, workspaces or configuration that is not yours.
- Do not probe, scan or load-test the platform without our written permission. Reporting a vulnerability you found in normal use is welcome and is not a breach — see [the security page](#).
- Do not attempt to bypass the safeguards described in section 8, including instructing the assistant to ignore them, or engineering documents or messages intended to override its instructions.
- Do not resell, sublicense or white-label the service without a written agreement, scrape it, or use it to build a competing product.
- Do not run automated traffic through a workspace at a volume that degrades service for others, or use the service primarily to generate AI output unrelated to supporting your own customers.
- Do not share one account between people, or leave accounts active for people who have left your team.

7. Content you upload

- You must have the right to upload what you upload, including any third-party material in it.
- Keep it accurate and current. The assistant repeats your documents faithfully, including their mistakes, and a stale price or policy will be quoted with confidence.
- Uploaded documents are indexed for retrieval. Do not upload anything you would not want quoted back to a customer.

8. What the platform enforces itself

Several limits do not depend on this policy being read. They are implemented in the software, which is why they hold even when a configuration asks otherwise:

- The assistant refuses to ask a customer for card numbers, CVV or security codes, PINs, passwords, one-time codes, bank account details or government identity numbers — recognised in English, French and Spanish — and refuses to store them as conversation state, whatever a workspace configuration requests.
- Phrases a workspace adds to its “never say” list are checked before a reply is sent. A reply containing one is withheld and the conversation goes to a person instead.
- A reply that promises something the platform cannot actually do is replaced before the customer sees it. Agreeing to an action records it as under way, never as done; only a step that genuinely completed can be reported as completed.

- Documents a customer sends are treated as unverified evidence, never as instructions and never as business knowledge: they are scanned for malware before anything reads them, are never added to the knowledge base automatically, and can never be used to confirm a payment or an identity. Infected files are deleted without being parsed.
- Channel activation and send readiness are checked according to the channel and its ownership. Health failures can block or disable the affected connection; one failed send does not automatically disable a platform-shared channel.
- When a usage limit is reached, conversations escalate to your team rather than the assistant failing quietly.

These are safeguards, not a substitute for judgement. They stop specific, nameable failures. They cannot stop a workspace pointed at the wrong job in the first place, which is what the rest of this policy is for.

9. What happens if this policy is broken

Our response is proportionate to the harm. Depending on the circumstances we may: contact you and ask you to fix it; disable a specific channel or feature; suspend the workspace, which stops it receiving and sending messages; or terminate the agreement under the [Terms of Service](#).

We act immediately and without prior notice where there is a serious risk of harm to people, to other customers, or to the platform, or where a channel provider or the law requires it. We will always tell you why, and we will lift a suspension once the cause is resolved.

10. Reporting abuse

If you have received a message from an assistant built on Tephlo that breaks this policy — a scam, harassment, a demand for card details — write to hello@tephlo.com. Include the phone number or account you were messaging, roughly when it happened, and what was said. We investigate reports about workspaces on this platform, and we can suspend one; we cannot act on messages that did not come through it.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.

Cookies and local storage

A short document, because there is very little to describe. No analytics, no advertising, no tracking — and a precise account of the few things that are stored.

Last updated 6 September 2026 · Revision Draft 1.1 — factual content refresh · Effective [EFFECTIVE DATE]

IN PLAIN LANGUAGE

The public Tephlo sites — marketing, documentation and this legal site — set **no tracking cookies of any kind**. There is no analytics script, no advertising pixel and no third-party script at all: the sites' security policy blocks them outright.

- One optional first-party cookie exists, on the pricing page, and only if you choose a different currency.
- The signed-in consoles set session cookies. Without them you cannot stay logged in, so they are strictly necessary.
- The embeddable chat widget stores a random id in the browser's local storage so a conversation survives a page reload. It is a pseudonymous conversation identifier, not a cookie.

1. The public sites

Visiting the marketing site, the documentation site or this legal site sets no cookies at all unless you take the action described in section 2. There is nothing to opt out of, because:

- no analytics product is installed — not a self-hosted one either;
- no advertising, remarketing or conversion pixels are present;
- application scripts and fonts are served by the site. Published logos and icons can load from configured API or asset-hosting origins allowed by the content security policy; those requests disclose ordinary connection data to the asset host;
- the public-site application does not implement browser fingerprinting or advertising tracking.

Our hosting provider processes ordinary connection data — IP address, request time, user agent — in server logs to deliver the page and protect against abuse, as any web host does. That is described in the [Privacy Policy](#), and the provider is on the [sub-processor list](#).

2. The one preference cookie

The pricing page shows prices in a market currency. If you switch markets, that choice is stored so the page keeps showing your currency next time.

SET ON THE MARKETING SITE ONLY

Name	What it holds	Set when	Lifetime
tephlo_market	A short market code such as a country or region identifier — nothing about you.	Only when you pick a market on the pricing page. Simply visiting the page sets nothing.	One year, first-party, SameSite=Lax.

An older cookie named `xephlo_market`, from before the product was renamed, is actively cleared when you set a market. Deleting either one simply returns the pricing page to its default currency. Your choice is never used to identify you, and the page never guesses your market from your IP address.

3. The consoles

The agent console and the operator console are private applications you sign into. They set strictly necessary cookies — without them there is no way to know you are signed in.

SET AFTER YOU SIGN IN

Cookie	Purpose	Lifetime
<code>__Host-dash_access / __Host-op_access</code>	Holds the short-lived access token that authorises each request.	Minutes — it is refreshed automatically while you work.
<code>__Host-dash_refresh / __Host-op_refresh</code>	Lets the console get a new access token without asking you to sign in again. Rotated on every use; replaying an old one revokes the whole session family.	Until it expires or you sign out.
<code>__Host-dash_mfa_enrollment / __Host-op_mfa_enrollment</code>	Marks that you are part-way through setting up multi-factor authentication, so the console knows where to send you. It contains no credential.	The enrolment step only.

All of these are `httpOnly` — JavaScript cannot read them, which is what stops a script on the page from stealing a session. In production they carry the `__Host-` prefix, meaning the browser refuses to accept them from any other host or path, and are marked `Secure` and `SameSite=Lax`.

The consoles also store display preferences locally: light or dark theme under `theme`, and sidebar state under `console:sidebar-collapsed`. The tenant console uses `kb-layout` for its knowledge-page layout. These preferences are not session credentials.

4. The embeddable chat widget

If a business puts the Tephlo chat widget on its own website, the widget stores a value in that site's local storage under `xw_visitor`: a random string the widget generates in your browser. It exists so that reloading the page does not start a new conversation from scratch.

- It is not a cookie, is not sent automatically with requests, and is readable only by that one site.
- It contains no name, email or account — it is random, and clearing site data replaces it with a new one.
- The conversation session itself is held on our servers for one hour of inactivity and then expires.
- Because it is storage on your device, the business embedding the widget is responsible for mentioning it in their own cookie or privacy notice.

5. WhatsApp and Telegram

Conversations on WhatsApp and Telegram happen inside those apps, and nothing in this document applies to them. What those apps store on your device is governed by their own policies.

6. Controlling what is stored

- **Browser settings.** Every browser can block or delete cookies and site data for a specific site. Nothing here resists that.
- **The preference cookie.** Delete it, or block cookies for the marketing site, and pricing simply shows its default currency.
- **The console cookies.** Blocking them prevents signing in — there is no alternative mechanism, which is what “strictly necessary” means.
- **The widget id.** Clearing site data for the site hosting the widget removes it.

Do Not Track and Global Privacy Control. We have nothing to switch off in response to these signals, because there is no tracking to begin with.

7. Why there is no cookie banner

Consent is required for storage that is not strictly necessary for a service the visitor asked for. On these sites there is none: no analytics, no advertising, no profiling. The currency cookie is set only by your own explicit choice, and the console cookies are necessary for a service you signed into. A banner asking permission for nothing would be theatre, and would train people to click through the ones that matter.

If that changes, this page changes first. If we ever add analytics, a consent mechanism will ship with it and this page will say so before it is switched on.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.

Security overview

What actually protects the data in a Tephlo workspace, described so a non-specialist can follow it and a specialist can check it — including the things we do not claim.

Last updated 6 September 2026 · Revision Draft 1.1 — factual content refresh · Effective [EFFECTIVE DATE]

IN PLAIN LANGUAGE

The short version: workspaces cannot see each other; sign-in uses short-lived tokens and supports multi-factor authentication; secrets and credentials are encrypted; sensitive values are stripped out of logs; administrative actions are recorded in trails that cannot be quietly edited; supported customer documents are scanned before parsing; and customer erasure removes the records it covers.

- These are implemented controls, not aspirations — each section describes something the software does.
- We hold no SOC 2 report and no ISO 27001 certificate, and no independent auditor has assessed any of this.
- Message content is not end-to-end encrypted: answering a question requires reading it.

Registrations on the public record

Two registrations stand behind the service, and they are held in two different names. Tephlo is the product; neither certificate carries that word. What follows is what the certificates say, no more, and where each record can be looked up.

TEPHLO AI VANGUARD SYSTEMS LTD

Certificate of incorporation of a private company limited by shares, issued by the Corporate Affairs Commission of Nigeria under the Companies and Allied Matters Act 2020. Reference RC 9893569, 28 September 2026.

What it establishes: That the company exists as a Nigerian private company limited by shares, with this registration number, from this date.

How to verify it: [Corporate Affairs Commission of Nigeria](#) — Search the Commission's public register by company name or RC number.

TEPHLO AI ENTERPRISES

Registration as a data controller / processor of major importance, issued by the Nigeria Data Protection Commission under the Section 44 of the Nigeria Data Protection Act 2023. Reference NDPC/DCP/15068, 28 September 2026 to 28 September 2027.

What it establishes: That the registration the Act requires of a data controller or processor of major importance has been filed and accepted for the period shown. It is not an audit of how data is handled.

How to verify it: [Nigeria Data Protection Commission](#) — The Commission publishes its register of registered controllers and processors; quote the registration ID.

The data-protection registration is held in the name TEPHLO AI ENTERPRISES, the registered business name (Corporate Affairs Commission business-name registration 9805539, 26 August 2026) under which the service was first offered. The company that contracts with customers under the [Terms of Service](#) is TEPHLO AI VANGUARD SYSTEMS LTD, incorporated on 28 September 2026.

What a registration is not. A registration records that a filing was made and accepted by the regulator. It is not a security audit, not a certification that data is handled correctly, and not an endorsement by the Commission or the government. The controls in the sections below are what protect the data; the registrations are the public record of who is answerable for them.

1. Keeping workspaces apart

Workspace data is scoped to its owning workspace. Authenticated routes check the caller's permissions, and channel and API routes bind requests to their authorised configuration. Supplying an identifier does not itself grant access to the records it names. Platform-wide configuration and audit records have separate operator controls.

A conversation is keyed by the workspace, the channel, the business identity that received the message and the customer — so the same person reaching two businesses, or reaching one business through two numbers, produces separate conversations that never merge.

Where several businesses share one platform number, which business a conversation belongs to is recorded using a salted one-way hash of the customer's identifier rather than the identifier itself.

2. Signing in

- Passwords are stored hashed, never in a readable form, and repeated failed attempts lock an account.
- Access tokens are short-lived — capped at fifteen minutes in production — so a leaked one is useful only briefly.
- Refresh tokens rotate on every use. If an old one is replayed, which is what a stolen token looks like, the entire session family is revoked rather than the attacker being handed a fresh token.
- Permission is re-checked against the database on every request, so disabling an account, changing a role or suspending a workspace takes effect immediately rather than when a token expires.
- Multi-factor authentication is enforced for platform operator accounts in production and available to workspace administrators; the authenticator secret is encrypted at rest.
- Console sessions live in host-only, `httpOnly` cookies that JavaScript cannot read, protected by an origin check on every state-changing request. The [cookie notice](#) names them.

3. Encryption

In transit: everything runs over HTTPS. The production configuration refuses to start if a provider endpoint is not HTTPS, if the connection to the cache is not encrypted, or if a managed database connection is downgraded to plaintext — a misconfiguration fails loudly instead of quietly running unencrypted.

At rest: the managed database and object storage apply their providers' own encryption at rest, which is their commitment rather than ours to guarantee. On top of that the platform encrypts specific things itself, so that a copy of the database alone does not hand over credentials:

- WhatsApp and Telegram credentials for every connected channel;
- multi-factor authentication secrets and invitation payloads;
- the bodies of operational emails, which are cleared once delivery finishes;
- uploaded knowledge documents while they wait to be indexed, and the text extracted from customer documents.

Encryption keys are held in the deployment's secret manager, and the channel key supports rotation: a new key becomes the encrypting key while the old one remains decrypt-only until everything has been re-encrypted.

4. Making sure a message really came from the channel

Inbound WhatsApp webhooks are verified with a cryptographic signature over the original request body, checked against the secret of the channel that owns the number. A forged webhook does not reach the pipeline.

Channel activation depends on the readiness checks for that channel. Health and credential failures can block delivery or disable an affected connection. The exact response depends on channel ownership and failure type; one workspace's send failure does not automatically disable a platform-shared channel.

5. What ends up in logs

Logs are for diagnosing faults, and they are one of the easiest places for personal data to leak. Email addresses, phone numbers, URLs, bearer tokens, signatures and stack traces are redacted at the logging boundary, so a value has to survive that filter to be recorded.

Operational analytics store categories and counts — how a conversation turn ended, how long a reply took, how many tokens it used — with no message content. Some records carry a conversation reference and are removed with that conversation. Where the platform needs to recognise a repeated question, it uses a hash of the normalised request rather than the text.

6. Audit trails

Actions that change security-relevant state are recorded with who did it, what they touched and the reason they gave: channel credential changes, routing changes, platform configuration, privacy exports and deletions, media operations, and workspace or account deletion.

Two of these trails are append-only at the database level: an attempt to modify or remove a row is rejected by the database itself, not merely by the application. Audit rows never carry message content, and privacy audit rows identify the person by a salted hash rather than by their phone number or user id — so the record that something was erased does not recreate what was erased.

7. Files a customer sends

Reading an untrusted file is one of the riskiest things a system can do, so the order of operations matters:

- **Scanned before anything opens it.** A file is checked by a malware scanner before any parser touches the bytes, because probing a file to identify it is already parsing it. If the scanner is unavailable, the file is not parsed at all — the check fails closed.
- **The bytes decide what it is.** The filename, extension and the type the provider claims are ignored. Executables, archives, macro-enabled documents, password-protected files and PDFs containing active content are rejected outright.
- **Extraction runs under hard limits** on pages, characters, rows and time, and spreadsheet formulas are treated as inert text.
- **Infected files are deleted without being parsed**, recorded as a security event, and can never be downloaded through the application. The customer gets a neutral message that names no scanner or signature.
- **Storage is private.** Files are stored under opaque keys the backend generates, partitioned per workspace, with no public access and no shareable links. Customer filenames never appear in a storage key.

8. The AI layer

The assistant is treated as a component that can be wrong or manipulated, and the guardrails around it are ordinary code rather than instructions to the model:

- **Customer content is handled as untrusted data.** The assistant is instructed to treat messages and documents separately from its rules, and code checks constrain supported actions. This does not guarantee immunity to prompt injection or other AI errors. Customer documents are never added to the knowledge base automatically, and are never used to confirm a payment or verify an identity.
- **Sensitive fields are refused.** The assistant will not ask for card numbers, security codes, PINs, passwords, one-time codes, bank details or government identity numbers, and will not store them as conversation state, whatever a configuration requests — recognised in English, French and Spanish.
- **Promises are checked.** A reply claiming an ability the platform does not have is replaced before sending; an action agreed to is recorded as under way and can only be reported as done once it genuinely completed.
- **Reuse is gated.** Answers containing email addresses, phone numbers, long card-like numbers or labelled codes are kept out of the answer-reuse cache, and text stored for optional quality review has those values masked first.
- **Customer memory has separate controls.** It is enabled by default in the software, can be switched off by a workspace and can be blocked by sensitivity controls. Newly extracted facts have a 180-day expiry and are filtered before storage; expired facts are excluded on read. Administrators can inspect and forget entries.

- **Smart Updates requires recorded consent.** Send checks consider the customer's opt-out, topic, frequency and channel eligibility. Each update includes a STOP instruction. Consent and send records are part of customer export and erasure.
- **Conversation working state is short-lived** — it expires thirty minutes after the last message, is kept separate from durable customer memory, and is re-validated against the same limits when read back.

Conversations are not used to train models. They are sent to the AI provider to generate a reply. The [sub-processor list](#) names that provider and says what reaches it.

9. Deletion and retention

Temporary state and retained records have different lifecycles. Read-time expiry stops expired memory being used. Retention workers can remove expired facts, inactive conversations the assistant handled, delivery records and duplicate-delivery protection records. Physical deletion depends on the relevant worker being enabled and running. Customer document processing attempts to remove raw files after text extraction by default; cleanup failures can require a later retry.

A workspace administrator can export or erase one customer's data. The erasure is serialised with message processing, so work accepted just before it either completes and is removed, or is suppressed afterwards — a reply cannot arrive after someone asked to be forgotten. The [Privacy Policy](#) lists exactly what it covers, including remembered facts, care cases, Smart Updates records, saved carts, potential-order requests and appointments. Operational records erased with those aggregates are not all separately exported.

Deleting a whole workspace removes everything it owns in a single transaction — either all of it goes or none of it does — and the files in object storage are scheduled for cleanup in that transaction. Removal requires the cleanup worker and storage provider to succeed; committing the request is not proof that every remote object has already gone. Platform audit records survive by design, and carry no message content.

10. Backups and recovery

The repository provides an encrypted database-backup procedure and retention settings. The active backup schedule, storage location and key management must be confirmed for the deployment; the presence of those scripts alone does not prove backups are running.

Restoring has a documented procedure and integrity checklist. A test restore into a disposable database does not establish recovery times for a production deployment. Operators must also reapply relevant erasures when restoring older data.

Because backups are not edited individually, data you delete leaves them as they age out and are pruned, not at the moment of deletion.

11. How changes reach production

- Every change goes through version control and automated checks: linting, type checking, an extensive test suite, dependency auditing, and database migration checks.

- Risky capabilities ship behind switches that default to off — document reading, image reading, transcription and web search all had to be deliberately enabled.
- The production configuration validates itself at startup and refuses to run with weak keys, an unencrypted datastore connection, disabled multi-factor enforcement, a non-HTTPS provider endpoint or unprotected metrics. A misconfigured deployment fails to start rather than serving traffic insecurely.
- Public surfaces carry strict security headers and a content security policy, request size limits, and rate limiting.

The [product documentation](#) describes how these appear in the console.

12. What we do not claim

Read this section as carefully as the rest. A security page that lists only strengths is marketing.

- **No certification.** No SOC 2 report, no ISO 27001 certificate, and no independent audit of the controls above. The registrations at the top of this page are statutory filings, not audits.
- **No uptime guarantee.** No service level agreement and no service credits.
- **No data residency guarantee.** Providers operate internationally and the AI provider may route a request onward. See the [sub-processor list](#).
- **No end-to-end encryption of messages.** The platform reads a message in order to answer it, and staff with production access could in principle see conversation content. Access is limited to people who need it to operate the platform.
- **Redaction is best-effort.** Masking of card-like numbers, codes and contact details is pattern matching. It catches the common cases; it is not a guarantee, which is why short retention rather than filtering is the real protection.
- **Duplicate delivery is possible in rare cases.** Messaging providers that offer no idempotency key cannot guarantee a message is delivered exactly once if a crash happens at exactly the wrong moment.
- **Scale is not certified.** Performance and isolation under very large load have not been independently verified.
- **Organisational controls are still being formalised.** Background checks, security training and formal policy review are marked as open items in the [DPA security annex](#) rather than claimed here.

13. Reporting a security problem

If you find a vulnerability, tell us at [\[SECURITY CONTACT\]](#) or hello@tephlo.com. Include enough detail to reproduce it, and please do not access, modify or retain anyone else's data while investigating.

Our side of it. We acknowledge reports, keep you informed while we fix the issue, and will not pursue anyone who reports a genuine problem in good faith. We do not run a paid bug bounty, so we will not pretend a reward is coming.

Contact

General questions about this document go to hello@tephlo.com. Questions about personal data, including requests from individuals, should go to the data protection contact at **[DATA PROTECTION CONTACT]**; until that address is published, the general address above reaches the same team. A request about your own data — access, erasure, correction, objection or a complaint — can also be sent through the [privacy request form](#), which puts it in front of the team as a tracked request rather than a message in a mailbox.

Postal address: **[REGISTERED COMPANY NAME]**, **[REGISTERED ADDRESS]**.