

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

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 2026-10-03T08:44:51.249Z. The pages are the authoritative text.

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]**.