

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]

TEPHLO AI ENTERPRISES

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

Tephlo · hello@tephlo.com

<https://legal.tephlo.com>

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

IN PLAIN LANGUAGE

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]**.