MAKING YOURSELF **CLEAR**

PRIVACY & SECURITY

Making Yourself Clear — Privacy & Security Summary

> *One-page summary for prospective customers' security and legal teams. Last reviewed: **2026-07-08**.*

Making Yourself Clear is an AI-powered leadership communication coaching platform operated by Executive Sound Board. Customers upload meeting transcripts; Making Yourself Clear analyses them and returns structured coaching feedback. The platform is multi-tenant: each customer organization holds its own coachees, coaches, and administrative controls.

This document summarises how transcript data is handled and answers the most common questions from security reviewers. A longer technical companion is available in `security-whitepaper.md`, and a detailed data-flow walkthrough with a diagram in `data-flow.md`.

Data flow at a glance

1. A coachee uploads a meeting transcript to their account over HTTPS.
2. Making Yourself Clear **redacts personally identifiable information** from the transcript before any LLM call (see *Redaction* below).
3. The redacted transcript is sent to the configured LLM provider. By default that is OpenAI, which does not train on API data and receives every request with `store=False` (disabling OpenAI-side conversation storage); standard provider retention (up to 30 days, for abuse-monitoring) applies. Customer organizations can configure their own endpoint (BYO LLM) so transcripts go through their own contract, not ours.
4. The structured coaching response is stored against the coachee's account and rendered in the app.
5. Each organization's data is isolated from other organizations via row-level tenancy keys and enforced at the API boundary.

What we collect and store

Data	Stored	Notes
Raw transcript text	Yes	Accessible only to the coachee, their assigned coach (if any), the org_admins of the coachee's organization, and authorized Making Yourself Clear vendor_admins.
Redacted transcript text	Yes	The exact text sent to the LLM. Visible to coachees and coaches from the run detail page via the "View what was sent" panel.
Redaction audit log + token mapping	Yes	Used to verify what was redacted; not used to substitute original values back into the AI output. Mapping can be disabled per-organization.
AI analysis output	Yes	Structured coaching JSON.
Account profile	Yes	Email, display name, role (vendor_admin / org_admin / coach / coachee), organization membership, authentication method.
Organization metadata	Yes	Name, status, member list, per-org privacy settings.
Administrative audit log	Yes	Every administrative action (role change, settings change, member invite) is recorded with actor, target, timestamp, and originating IP / User-Agent. Append-only.
Operational logs	Yes	Standard application logs (timestamps, endpoints, status codes). Forwarded from Railway to Better Stack with a 3-day retention on the current plan. PII is not deliberately logged; the redaction audit log records replacement tokens, not original values.
Encrypted BYO LLM API keys	Yes	When an organization configures a BYO LLM endpoint, their API key is encrypted at rest with AES-256-GCM under the deployment root key (KMS-rooted on SaaS). Plaintext never appears in logs, API responses, or the audit log.

We do **not** collect or store: the contents of inboxes, calendars, files outside the transcripts you upload, behavioural analytics beyond what is required for security, or marketing/advertising identifiers.

Redaction (pre-LLM)

Making Yourself Clear runs a redaction layer (Microsoft Presidio plus custom recognizers) on every transcript before any LLM call. Items redacted by default:

- Names of people **not present in the meeting** (third parties mentioned in conversation).
- Email addresses, phone numbers, physical addresses.

- Government IDs: US Social Security numbers, passport numbers, driver's licence numbers, medical licence numbers.
- Credit card numbers and partial card references.
- Employee / staff IDs (EMP-####, etc.).
- Dates of birth.
- API keys, tokens, webhook secrets (OpenAI, Stripe, generic labelled secrets).
- IP addresses, URLs.
- Locations (in Conservative aggressiveness mode), nationalities / religions / political groups.

Items **intentionally preserved** (with reasoning):

- **Names of the meeting participants (speakers).** The LLM needs to refer to speakers by name to produce coherent coaching feedback. The coachee's own name is already in their Making Yourself Clear account; preserving speaker names does not expand exposure. The full name and individual name tokens are protected by the do-not-redact allowlist.
- **Organization names.** Off by default. Can be enabled per-deployment by the vendor admin, and per-organization by an org admin (override).

Per-organization overrides. Each customer organization can override the global defaults for five settings (enabled / aggressiveness / reversible-mapping / redact-organization-names / share-original-with-coach). Settings left unset inherit the vendor defaults. Org admins manage this in the in-app Settings page.

In-app verification. Coachees and coaches can open the "View what was sent to the LLM" panel on any analysis to see the exact redacted text.

Customer privacy options (privacy ladder)

Making Yourself Clear supports three deployment postures depending on a customer's trust requirements:

Tier	What you get	When it applies
Tier 0 — Default SaaS	Hosted by Making Yourself Clear, with the default redaction pipeline, application-layer at-rest encryption (per-organization key, KMS-rooted), and a no-training LLM provider (`store=False`). Each organization's data is logically isolated.	Most customers. The platform default.
Tier 1 — BYO LLM endpoint	The customer configures their own LLM endpoint (Azure OpenAI deployment, enterprise OpenAI key, OpenAI-compatible proxy, or Anthropic-direct) at the organization level. Transcripts never touch Making Yourself Clear's LLM provider contract — only the customer's.	When the customer's LLM provider relationship needs to govern the data. Configured by the org_admin in-app.
Tier 2 — Self-hosted	The customer runs the Making Yourself Clear stack (frontend, backend, worker, database) in their own VPC. Transcripts never leave the customer's infrastructure.	Customers who require true vendor-zero-knowledge (e.g. regulated industries, government). Generally available via a turnkey self-host installer.

A Tier 3 (fully local LLM, no cloud LLM at all) is currently out of scope; small open-weight models do not yet match frontier-model coaching quality.

AI provider posture

- **OpenAI** (default routing): OpenAI does not train on data submitted via its API (standard API terms). All requests are sent with `store=False`, which disables OpenAI-side storage of the conversation. Standard OpenAI API retention of up to 30 days (for abuse-monitoring) still applies. A Zero Data Retention (ZDR) agreement with OpenAI has been **applied for and is pending** — it is not yet in effect, so we do not claim zero retention today.
- **Anthropic** (supported for Claude models): The application's LLM router selects Anthropic when the configured model name has a `claude-` prefix. Anthropic-direct traffic is also subject to the provider's no-training default. Customers who want to route exclusively to Anthropic can do so via the per-org BYO LLM configuration.
- **BYO LLM endpoints** (per-organization): An org_admin can configure their own LLM endpoint and API key. Retention is then governed by the customer's own contract with their LLM provider. The customer's API key is encrypted at rest and never returned in plaintext through the UI.

Making Yourself Clear does not use customer transcripts to train models, evaluate prompts, or improve internal systems without explicit written consent.

Storage, encryption, access controls

Control	Posture
Encryption in transit	TLS 1.2+ for all client ↔ server and server ↔ LLM-provider traffic.
Encryption at rest	**Application-layer AES-256-GCM** over all customer content (raw + redacted transcripts, token mappings, analysis output, baseline summaries) under a **per-organization data-encryption key**, plus all at-rest secrets (BYO LLM keys, OIDC/OAuth client secrets, etc.). On SaaS the root key is held in **AWS KMS**, adding decrypt audit + key revocation / crypto-shred. Standard cloud-managed AES-256 volume encryption (Railway managed Postgres on GCP), including backups, applies beneath.
Authentication	Google OAuth, Microsoft OAuth, or email + password with mandatory email verification. Users authenticating via OAuth inherit MFA from their identity provider. Elevated roles (coach, org_admin, vendor_admin) require an OAuth credential, so all elevated-role users inherit MFA from their IdP. Email + password coachees do not yet have application-level MFA — planned.
Role-based access	Four roles: **vendor_admin** (Making Yourself Clear operator), **org_admin** (administers their organization's members and settings), **coach** (sees only their assigned coachees), **coachee** (sees only their own data). Permissions are enforced at the API boundary on every request.
Multi-tenant isolation	Each domain row (transcript, run, baseline pack, experiment) carries an `organization_id`. Cross-organization access is blocked at the query layer. Cross-organization probe attempts return 404 (not 403) to prevent enumeration.
Admin audit log	Every administrative action (role change, redaction settings change, BYO LLM change, member invite/removal) is recorded to an append-only `admin_audit_log` table, capturing actor identity, action, target, timestamp, and originating IP / User-Agent. The BYO LLM API key value is never written to the audit log — only the actions "set" / "rotated" / "cleared" are recorded.
Background workers	Run in the same trust boundary as the application server.

Shared responsibility model

Responsibility	Making Yourself Clear	Customer
Securing the application code and runtime	✅	—
Patching the OS and dependencies in our infrastructure	✅	—
Encrypting data in transit and at rest	✅	—
Pre-LLM PII redaction	✅ (configurable)	Tuning aggressiveness and overrides for their org
Enforcing role permissions and multi-tenant isolation	✅	—
Sub-processor selection and management	✅	Reviewing the published sub-processor list
Choice of LLM provider for the organization	Default (OpenAI)	Optional BYO LLM endpoint configuration
Managing their own LLM provider contract (if BYO)	—	✅
Account access controls (who in the org has what role)	—	✅ (org_admin)
Endpoint security on coachee devices	—	✅
Acceptable-use policy (what is uploaded)	—	✅

Data retention and deletion

- **Self-service deletion:** Coachees can delete any transcript and its analyses from the run detail page. Deletion is immediate from the user's view; backups age out per the backup policy below.
- **Account deletion:** Coachees can self-delete individual transcripts and runs. Full account deletion is by request to security@makingyourselfclear.com and is processed within 30 days.
- **Organization removal:** When a member is removed from an organization, their domain rows are unscoped from the organization (the organization_id is cleared) but the rows themselves are preserved on the coachee's account. The org_admin can also be removed by a vendor_admin; the same row-scoping applies.

- **Backup retention:** Railway managed Postgres provides daily automated backups with point-in-time recovery on the current plan. Restore procedures are tested at major releases.
- **Operational logs:** 3 days via Better Stack's retention policy on the current plan; logs older than that are deleted automatically.
- **Departure of customer:** On request at termination, customer data is returned in a machine-readable format; otherwise it is deleted within 90 days of the end of the subscription.

Sub-processors

Making Yourself Clear uses the following sub-processors to deliver the service. **The list is kept up to date; material changes are notified per the DPA.**

Sub-processor	Purpose	Data location
OpenAI	LLM inference (default routing)	US, OpenAI's default routing
Anthropic	LLM inference for Claude models (via the model-prefix router; routed only when a Claude model is configured)	US
Railway	Application hosting (backend, worker, frontend, managed Postgres). Built on Google Cloud Platform infrastructure.	US West (California)
AWS KMS	Root-key custody for at-rest encryption (wraps the per-organization data-encryption keys; never receives customer content)	US East (Ohio)
SendGrid	Transactional email (verification, reset, invite)	United States
Better Stack	Centralised application logs and uptime monitoring	European Union
Sentry	Error tracking	United States
GitHub	Source control and dependency / secret scanning	United States
Google	OAuth identity provider (when chosen by user)	n/a (auth handshake only)
Microsoft	OAuth identity provider (when chosen by user)	n/a (auth handshake only)

When a customer configures a BYO LLM endpoint, the customer's chosen LLM provider becomes a customer-side sub-processor governed by the customer's own contract rather than Making Yourself Clear's.

Certifications and audits

- **SOC 2 Type II:** In progress. Engaging an attestation firm (Drata / Vanta evaluation) in 2026; Type I expected within 6 months of engagement, Type II to follow. Until issued, we share the controls posture documented here and in `security-whitepaper.md`.
- **Independent penetration test:** Planned. Scheduling a focused web-application engagement (NCC Group / Trail of Bits / Bishop Fox or equivalent). Reports are shared under NDA with prospective customers.
- **ISO 27001:** Not yet pursuing.
- **GDPR:** Making Yourself Clear processes personal data on behalf of customers as a Data Processor. A Data Processing Agreement is available — see `dpa-template.md`. The template requires Ontario-licensed legal review before signing.
- **HIPAA:** Making Yourself Clear is **not** intended for use with Protected Health Information. Customers should not upload PHI; we do not offer a Business Associate Agreement at this time.
- **CCPA / state privacy laws:** Making Yourself Clear acts as a Service Provider under CCPA/CPRA for its customers. We process Personal Information only on documented instructions from the customer and do not sell or share Personal Information.

What we don't do

For transparency, an explicit list of things Making Yourself Clear does **not** do today. These are limits that prospective customers should understand before relying on the platform for their highest-sensitivity workflows.

- **We do not offer true vendor-zero-knowledge in the default SaaS tier.** Our worker reads transcript plaintext momentarily during analysis. Customers who need cryptographic isolation from the vendor should use Tier 2 (self-host).
- **We do not currently support *customer*-managed encryption keys (CMEK / BYOK) at rest.** Customer content is encrypted at rest with application-layer AES-256-GCM under per-organization keys, and on SaaS the root key is held by us in AWS KMS — but a key that the *customer* holds and can revoke is not yet offered. CMEK/BYOK is on the roadmap; customers who need to hold the keys themselves today use Tier 2 self-host.
- **We do not train any models on customer data.** Not our own models, not third-party models. `store=False` is set on every OpenAI call; equivalent posture is documented for any BYO LLM endpoint configured by a customer.
- **We do not sell, share, or otherwise monetize customer data** to any third party.
- **We do not currently support data residency selection** beyond our default region (US West). Customers with strict residency requirements should engage with us about Tier 2 self-host.

- **We do not have a Business Associate Agreement** for HIPAA-covered data. Do not upload PHI.

Incident response

- **Detection:** Application logs are forwarded from Railway to Better Stack for centralised storage and querying. Errors are tracked via Sentry with privacy-conservative defaults (no PII, no request bodies, no stack-frame local values). Uptime monitoring via Better Stack pings `/health` every three minutes, with email alerts to the on-call founder. No dedicated SIEM at present; SIEM is planned alongside the SOC 2 Type II trajectory.
- **Notification:** Making Yourself Clear will notify affected customers of a confirmed personal-data breach within 72 hours of becoming aware of it, in line with GDPR Article 33.
- **Vulnerability disclosure:** Security researchers can report vulnerabilities to security@makingyourselfclear.com. We commit to acknowledging reports within 3 business days.
- **Customer notification of sub-processor changes:** Material sub-processor changes are notified per the DPA, with a defined objection window.

Contact

- **Security questions:** security@makingyourselfclear.com
- **Privacy / data-protection questions:** privacy@makingyourselfclear.com
- **Procurement / DPA negotiation:** chris.iskander@makingyourselfclear.com

Document control

- Version: **2.1** (2026-07-08) — reflects the current architecture: application-layer at-rest encryption + AWS KMS rooting, the current LLM-provider posture, and Tier 2 self-host (generally available)
- Owner: Chris Iskander, Executive Coach and Founder, Executive Sound Board
- Last reviewed: 2026-07-08
- Next scheduled review: Annual review, or on material change to the data flow