


A private AI assistant that runs on infrastructure dedicated to one person. This brief states what is true today and clearly marks what is planned. It is written so a security reviewer can approve individual use without encountering a claim that live inspection would contradict.

1. What PurposeMapped is

PurposeMapped provisions a **dedicated single-tenant instance** (one AWS EC2 virtual server per user) running an AI assistant with persistent memory, a workspace of files, and three agents (Chat, General, Dreaming). It is **not** a multi-tenant SaaS app: your memory and files live on your own instance, not in a shared store and not inside OpenAI's, Anthropic's, Microsoft's, or Google's product cloud.

2. Data-flow



Dashed = control/health metadata (no user content). Solid = user content within the trust bound

User content stays on the dedicated instance. Model API calls carry the request to the provider under commercial API terms. Only sanitized operational metadata (component, error class, counts — never user text) flows to the PurposeMapped control plane for provisioning, updates, and health.

3. What is stored, and where

Data	Location	Encryption
Assistant memory, workspace files, PurposeMapped profile, Companion Charter	Your dedicated EC2 instance (EBS volume)	Encrypted at rest (EBS)
Account record (email, hashed password, plan, subscription status)	PurposeMapped control-plane database	At rest on host; passwords hashed
BYOK provider API keys (Tier 1)	Encrypted (Fernet) during provisioning, then deleted from the control plane after the instance is configured	Envelope-encrypted; not retained
Chat transport	Teams / Slack (your tenant) or a private Matrix room	Per channel; Matrix user-room encryption default-on (planned, see §8)

4. Model providers & their API terms

PurposeMapped calls frontier models through **commercial API terms, not consumer-app terms**. This matters: consumer chat apps have had murkier training defaults and, in one 2025 case, a court-ordered retention obligation. API traffic was excluded from that and is governed differently.

Provider	Training on your data?	Retention
OpenAI API	No (API/Enterprise inputs & outputs not used for training by default)	~30-day abuse-monitoring window; Zero-Data-Retention available on approval for eligible endpoints
Anthropic API	No (API-with-your-key is non-training)	Zero-Data-Retention available for eligible organizations (as of 2026-06-09)
Google Vertex AI	No training by default	Contractual ZDR-equivalent for enterprise via DPA amendment

Honest scope note: At the *model-API layer*, PurposeMapped's retention/training posture is **equal to** these vendors' enterprise/API terms — we say "equal," not "better." PurposeMapped's distinct privacy value is **architectural**: your memory and files live on an instance dedicated to you rather than in a vendor's product cloud, and there is no central PurposeMapped content store aggregating across users.

Sources: OpenAI business-data & enterprise-privacy pages (2026); Anthropic API terms / ZDR eligibility (2026-06-09); Google Vertex AI zero-data-retention documentation. Full citations available on request.

5. Subprocessors

Subprocessor	Purpose
Amazon Web Services (AWS)	Compute, storage (EC2/EBS), network for the dedicated instance and control plane
Anthropic, OpenAI, Google	Model inference under API terms (whichever the plan/role uses)
Stripe	Payment processing (billing metadata only; no assistant content)
Google Workspace / Gmail (optional)	Email & calendar integration, only if the user connects it

6. Security controls (today)

- **Dedicated instance per user** — no multi-tenant data commingling.
- **Encryption at rest** — EBS-encrypted volumes.
- **Closed public SSH** — administrative SSH to production is **closed to the public internet** (removed 2026-09-30); access is via private network / AWS SSM Session Manager break-glass only.
- **Recoverable off-host backups** — nightly encrypted backups to independent storage; weekly instance snapshots.
- **Managed updates** — platform fixes delivered through a signed updater channel; versions pinnable for company deployments.

7. Operator-access policy

Honesty is the point of this section. Today, PurposeMapped operations **technically can** reach instances for provisioning and support (administrative key / AWS SSM). The product promise is **policy + technical separation**, stated plainly:

- PurposeMapped staff **do not read your conversations or files** as a matter of policy. Support is consent-based.
- The chat surface is the user's; there is no standing "operator" account sitting in your conversation.
- Support increasingly runs **without any human reading content**: the assistant emits a sanitized failure signature, PurposeMapped ships a fix as a skill/update, and the instance applies it.

Planned, not yet live (do not treat as in-place): a user-visible **Access Ledger** that logs every administrative touch (session, support grant, config/skill push) and shows it in your dashboard is **targeted for Q4 2026**. Technically-enforced, logged, time-boxed consent for any human access is being implemented alongside it. Until then, the honest statement is "policy-based, consent-based, not yet independently user-auditable."

8. Encryption & transport posture (status)

- At rest: **EBS encryption in place**.
- Matrix user-room encryption default-on and stale federation advertising removed: **in progress**. Companies may instead use their own Teams/Slack tenant as the chat surface.

9. Retention, deletion & export

- **Export Everything** — a self-serve export of your memory, PurposeMapped profile, Companion Charter, and workspace files (portability = proof of ownership). We do not keep a separate store of your conversation transcripts — what persists is the memory and files on your own instance.
- **Deletion** — on cancellation the instance is suspended and, after the defined lapse window, decommissioned; backups age out on the documented schedule.
- Raw assessment answers and the full report are treated as sensitive user content and are never shared with marketing systems.

10. Incident contact

Security questions and incident reports: **support@purposemapped.com** (mark "Security"). A security reviewer may request the current subprocessor list, the data-processing terms, and a point-in-time control attestation.