MARYLINK

June 2026 edition · for CIOs, CISOs and compliance leads. This document summarizes Marylink's security posture. Contractual items (DPA, list of subprocessors, incident procedure) are shared on request and tailored to your case.

1. What Marylink is — and isn't

Marylink is a layer that **capitalizes on AI practices**. It does not replace your assistants (ChatGPT, Claude, Copilot, Mistral, internal agents): it structures, protects and makes reversible what your teams build with them. Direct security consequence: **Marylink does not create a new model-exfiltration surface** — it keeps the asset (the practice) on the organization's side, governed and exportable.

2. Data & confidentiality

- **Your data is never used to train third-party models.** This is a principle, reflected in our terms and in our vendor integrations (enterprise terms).
- **Strict partitioning per space and per client.** One client's data is never mixed with another's.
- Only the **context required** for a request is sent to the model, under the vendor's enterprise terms.

3. Hosting & encryption

Topic	Measure
Location	**European** infrastructure (servers in Germany), operated by Marylink. No data outside the EU in normal operation.
Encryption in transit	**TLS 1.3** end to end (HTTPS everywhere).
Backups	Automatic, **encrypted (restic)**, **daily off-site copy**, rolling retention (days / weeks / months).
Restore	Restore drills **run regularly**.

4. AI models

Marylink relies on **leading models, operated from Europe**. Only the context strictly necessary for a request is sent to the vendor, **under its enterprise terms**, and **never** to train third-party models. You keep the choice of model: switching vendor does not lose you the asset.

5. Permissions, access & governance

- Access organized by **space, role (champion, moderator, expert, member), group and validation step**.
- **Human** accountability preserved: a practice becomes "validated" because a reviewer endorses it, not automatically.
- Internal AI usage is **framed without being blocked**: shared practices, validations, traceability.

6. Traceability & auditability

- Who **published, validated, modified** what: everything is **logged and versioned**.
- Practices are **logged, versioned, reviewed and exportable**.
- A **security audit is possible before deployment**.

7. Reversibility & export — no proprietary lock-in

- **Native PDF export** of your practices.

- Access via the **MCP** standard: your own systems and agents can query your practices.
- On exit: **full export**, then deletion of data and backups per the schedule agreed in the contract. **You leave with your asset.**

8. GDPR & AI Act

- Product **designed to ease your GDPR and AI Act requirements** (practice registry, traceability, EU location, human accountability).
- **DPA** (data processing agreement) and **list of subprocessors** — with their location and transfer guarantees — shared on request.
- **Deletion procedure** on request: full export then erasure per the contractual schedule.

9. Continuity & incident management

- Encrypted backups + daily off-site copy + tested restore (see §3).
- **Documented incident-management procedure**, with client notification within the agreed timeframe. Details shared in a security review.

10. Transparency on certifications

We **do not claim any certification we have not obtained**. We document our measures and **support your assessments** (security questionnaire, architecture review, pre-deployment audit).

Talk to a founder

Executives, CIOs, legal: we discuss your trust and compliance requirements directly, on your case. **contact@marylink.io** · marylink.io

Informational, non-contractual document. The applicable commitments are those of the contract and its data processing annex (DPA).