

Data Handling & Security Brief

How ML Mind minimizes access while enabling safe AI savings.

Data principle

Only collect the minimum signals required for the requested deployment level.

Area	What to review
Logs-only	Request id, model, latency, tokens, status, error type and cost estimate.
Pre-model	Retrieved chunks, token budget and policy constraints when context control is enabled.
Gateway	Input, output, verification result and fallback decision when full control is enabled.
Self-hosted	GPU utilization, memory, queue, OOM and serving metrics.

Trust controls

Customer data is not used for model training. Enterprise deployments can be reviewed by privacy and security teams.

ML Mind - Safe AI Savings Control Plane - mlmind.cloud