

The engine runs on your data.

Here is everything that happens to it.

An AI-operated engine inside your pipeline is a fair thing to have questions about, so this page answers them before they are asked: what the engine touches, where it goes, what the AI layer keeps, and what you control. **Every claim here is verifiable**, and I will gladly walk your security or IT lead through any of it.

STANDING RULE

Your data works for you alone.

01 WHERE YOUR DATA LIVES

Your CRM stays the system of record. The engine works through access you grant on the stack you already run: it reads contacts, companies, threads, and engagement and writes sends, tasks, and outcomes back. Working files (the cadence board, the send ledger, the playbook) live in the engagement folder, which is yours at handoff. No shadow database, no bulk exports; the API surface is create and update only, so the engine cannot delete a record.

02 WHAT THE AI LAYER KEEPS

The engine runs on Claude, built by Anthropic. Prospect records pass through a working session while the engine researches, drafts, and verifies; they are encrypted in transit (TLS 1.2+) and at rest (AES-256). **The account behind your engagement has model training disabled:** nothing from your sessions is used to train AI models, and Anthropic deletes session data on a 30-day cycle. The policies are public; check them at privacy.claude.com and trust.anthropic.com.

03 IF YOUR POLICY NEEDS MORE

The engine is portable by design. If your posture requires enterprise AI terms, your engagement runs inside **your own Claude for Work organization** instead: your agreement with Anthropic, no training under its commercial terms, admin visibility, custom retention, a signed DPA, and SOC 2 Type II, ISO 27001, and ISO/IEC 42001 attestations behind it. Same engine, your tenancy.

04 WHAT I NEVER DO

Your data is never sold, never shared, and never shown to another client. It never seeds another engagement's targeting, research, or copy. Nothing reaches beyond the scopes you granted. One client per workspace, *walls between them*.

THE SHORT VERSION

trained on?	no. training disabled
sold or shared?	never, to anyone
system of record	your CRM, always
encryption	TLS 1.2+ / AES-256
AI-side retention	deleted on a 30-day cycle
who clicks send	a human you named, by count
revoke access?	any time, immediately
at handoff	files to you, access cut

05 GUARDRAILS YOUR PROSPECTS FEEL

Nothing sends without a named human approving the batch by explicit count, and every guard is built so a failure lands on **no send, never a bad one**. A genuine reply freezes that account instantly. An opt-out is honored permanently, on every channel. Your domain's reputation is treated as working capital.

06 WHAT YOU CONTROL

Access is minimum-scope and revocable at any moment, effective immediately. Every send lands in a ledger with its verified sent artifact, auditable on request. You own the research, the copy, and every output.

07 AT EXIT

The playbook, runbook, ledgers, and state files transfer to you; they were built in your folder to begin with. My access gets revoked on your side, so you never have to take my word that it happened. What stays is a working engine and your people trained to run it.

VERIFY IT YOURSELF: [TRUST.ANTHROPIC.COM](https://trust.anthropic.com) /
[PRIVACY.CLAUDE.COM](https://privacy.claude.com)
SECURITY REVIEWS WELCOME · [DAX@DAX.IO](mailto:dax@dax.io)