rereason.ai

Responsible Reasoning

Frontier AI reasoning on your most confidential data — without the data ever leaving your company.

Company overview · Vienna, 2026 · contact@rereason.ai · rereason.ai

Enterprises cannot use frontier AI on their most valuable data

Formulas, deal terms, pre-filing patents, process know-how: the decisions that matter most run on data that must not leave the company.

8.5%

of employee prompts to public LLMs already contain sensitive company data

90%+

of organizations see employees use AI tools — mostly unapproved “shadow AI”

0

legal protection left: a trade secret disclosed to a third party can lose its status

Why “we have a contract” is not enough

- Zero-retention clauses are promises, not proofs — a breach is invisible
- US CLOUD Act: extraterritorial access to data held by US providers
- Insider risk at the provider is outside the customer's control
- Pre-filing patents, M&A, formulas: irreversible damage if leaked once

Blocking does not work: employees route around weak AI

A policy is only as strong as the sanctioned tool. If the approved model cannot solve the problem, the employee opens a browser tab.

The employee's choice, every day

Sanctioned tool: local LLM

Underpowered

Private. Compliant. Descriptive answers, misses the hard part.

Browser tab: public frontier LLM

Where the work gets done

Solves the problem in one shot. Confidential data pasted straight into a third-party cloud.

What the data says

90%+ of organizations see employees use AI tools — most of it unapproved “shadow AI”

~40% of firms buy an official AI subscription at all — the rest run on personal accounts

8.5% of prompts to public LLMs already contain sensitive company data

You cannot police your way out of this. The only policy that holds is making the sanctioned tool the best tool — frontier reasoning, inside your walls. That is what rereason.ai delivers.

Our approach: semantic homomorphic transformation

Transform, don't block. A local model converts confidential context into a semantically faithful but generic representation. The frontier model reasons on it — and never sees the real data.

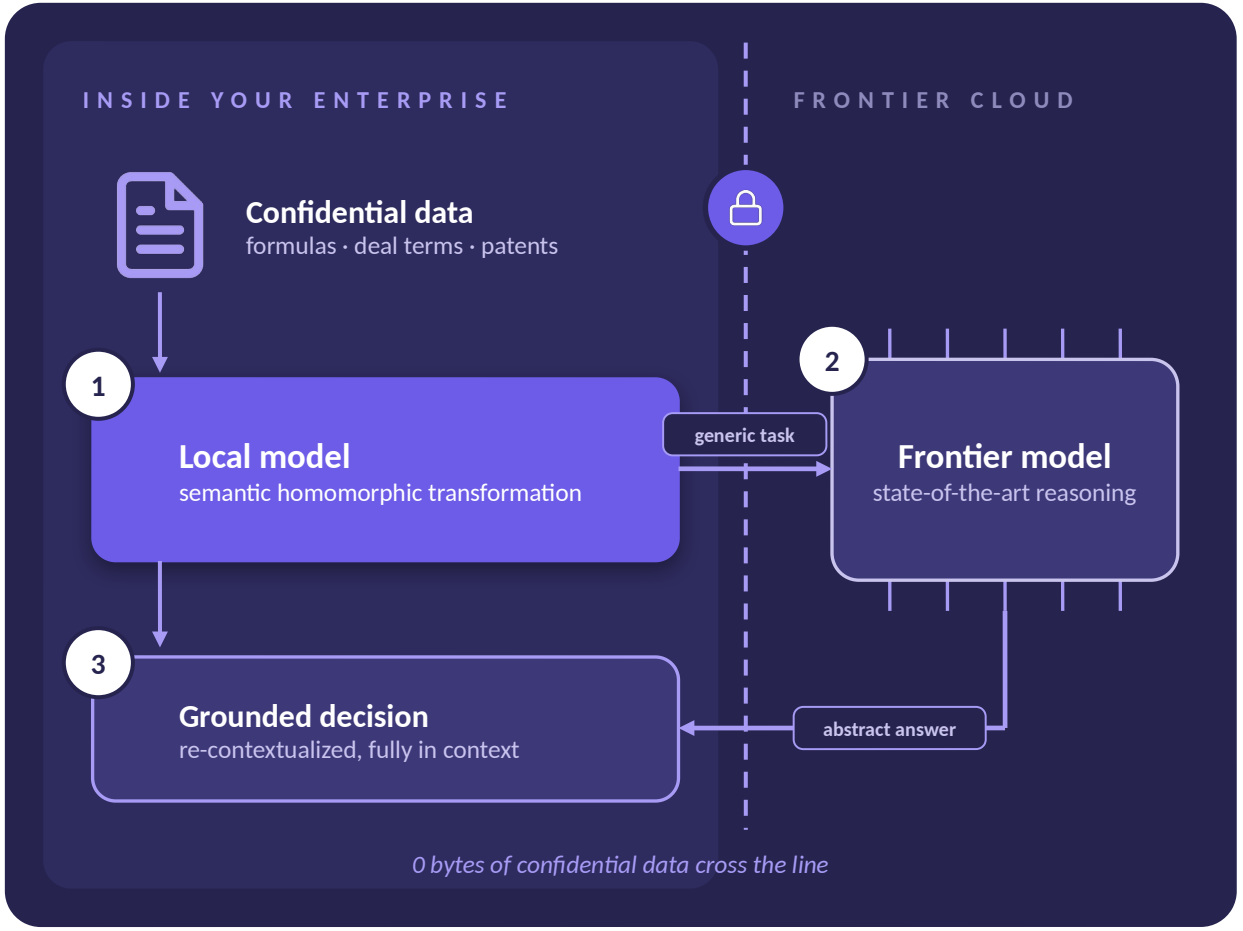

1 Abstract — inside your walls

Confidential data and business context are transformed locally into a generic, semantically faithful task.

2 Reason — at the frontier

The world's best models reason on the abstract task and return an abstract recommendation.

3 Re-contextualize — locally

The answer is grounded back into the real data, on-premise: a decision, fully in context.

Why "homomorphic": like homomorphic encryption, the transformation preserves structure — reason(transform(x)) ≈ transform(reason(x)) — but on semantics rather than ciphertext, at practical speed instead of cryptographic cost. Steps 1 and 3 never leave the enterprise boundary.

Working prototype: one real audit task, local LLM vs. rereason.ai

Same 16 KB confidential file, same task, both fully on-premise. Built in month 1 — measured, not promised.

Local LLM alone

Qwen 35B, on-premise

- Accurate but purely descriptive summary
- Zero issues found
- Nothing actionable
- Private — yet underpowered

rereason.ai prototype

semantic homomorphic transformation + frontier reasoning, on-premise

- Full structured audit: 11 severity-ranked findings, each with evidence
- Pass/fail across 7 dimensions
- Top-5 prioritized fixes
- Private — and frontier-grade. The data never left the machine.

	Local LLM	rereason.ai
Data stays local	Yes	**Yes**
Issues found	0	**11, severity-ranked**
Actionable output	No	**Yes — prioritized fixes**
Reasoning quality	Descriptive	**Frontier**

Team: research excellence meets commercial execution

David Saal

CEO & Co-Founder

- 10+ years consulting & IT product management
- Startup experience — from concept to delivery; at the moment building for IP-heavy companies
- Owns commercialization: design partners, pilots, funding

Prof. Michael Feischl

Head of Research & Co-Founder

- Professor of Mathematics, TU Wien
- 50+ peer-reviewed publications, multiple competitive research grants
- Leads a 10+ person research group — first-pick hiring pipeline of PhD-level talent for rereason.ai

TU

Research-anchored, clean IP by design. The scientific lead is a TU Wien professor; the TU Wien group is our talent pipeline and academic sounding board. All core technology is developed in-company with full IP assignment — no open university claims, no spin-off licence to negotiate. rereason.ai is in formation in Vienna.

Vision

The standard interface for the confidential use of frontier AI in Europe.

Every European enterprise should be able to use the world's best AI reasoning on its most sensitive decisions — with guarantees, not promises.

Design-partner program open for IP-heavy enterprises · Investors: contact@rereason.ai

David Saal · CEO | Prof. Michael Feischl · Head of Research