

Corentis Shield

AI checkpoint for regulated workflows

Investors and strategic funders

Corentis Investor Overview

A strategic overview of AI checkpoint infrastructure for regulated workflows.

AI needs a checkpoint before it acts. Corentis provides it.

A warm, commercial overview of the Corentis opportunity, the first wedge and the path from validation to regulated AI infrastructure.

Generated April 2026

For discussion and pilot exploration only

Overview

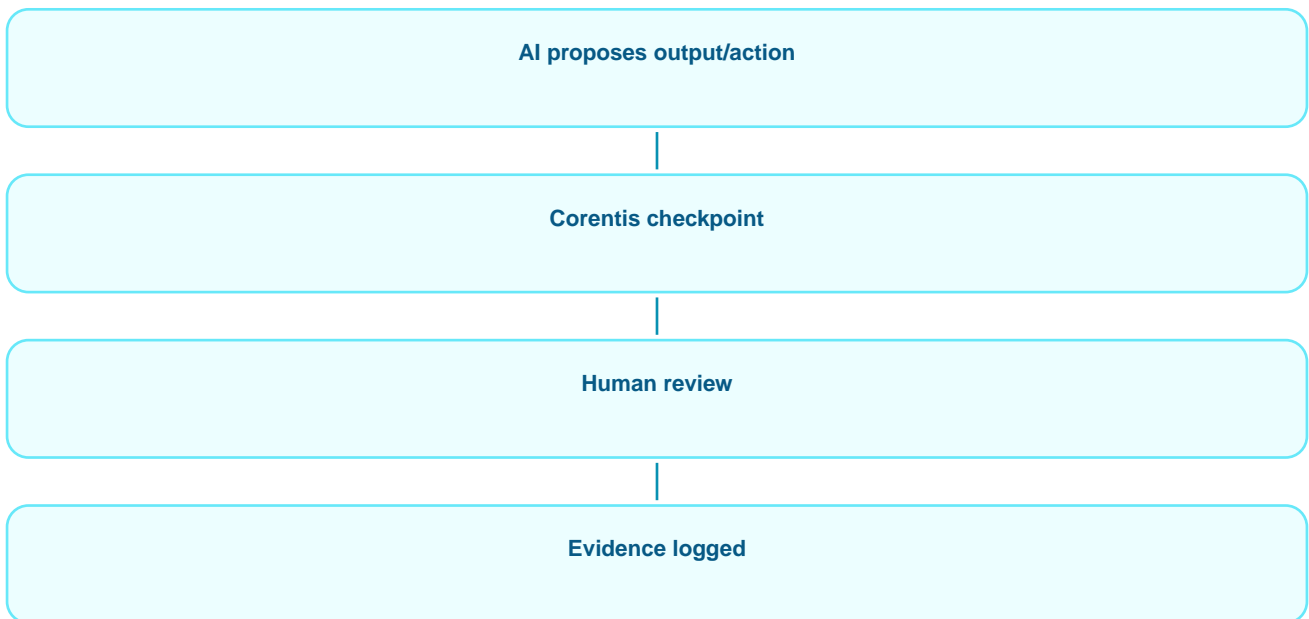
Core position

Corentis Shield is an AI checkpoint for regulated workflows.

AI needs a checkpoint before it acts. Corentis provides it. Corentis Shield is designed to help teams check AI outputs before they reach customers, teams or live systems.

VISUAL SUMMARY

Checkpoint flow



EVALUATION SHAPE

Baseline vs checkpoint

Baseline

AI proposes output or action without a runtime checkpoint. Review points and evidence gaps are assessed afterwards.

Checkpointed

AI proposes output or action. Corentis checks controls, pauses risky items, routes human review and records evidence before action.

A fast-growing market moment

Corentis sits at the intersection of three fast-growing markets: AI assurance, AI governance, and AI-enabled RegTech. The UK AI assurance market alone was valued at approximately £1.01bn GVA in 2024 and could reach £18.8bn by 2035 if adoption barriers are addressed. On that trajectory, the UK market would grow from roughly £1.7bn in 2026 to around £5.0bn by 2030 - nearly tripling in four years.

AI needs a checkpoint before it acts

AI is moving from drafting text to proposing actions, decisions and workflow steps. That shift creates a new infrastructure need: regulated organisations need a control point before AI-generated actions reach customers, case files or live operations.

The shift from drafting to acting

The first wave of enterprise AI helped teams write, summarise and search. The next wave is agentic: AI will suggest next best actions, draft customer responses, update records and trigger workflow steps. The question is no longer only what the model said. It is whether the action should move forward at all.

Why the first wedge matters

Corentis starts with financial-services complaints and vulnerable-customer workflows because they are high-consequence, evidence-heavy and deeply human. These workflows combine customer harm, regulatory sensitivity, operational pressure and a clear need for review before sensitive action.

A human moment

Example context: a customer discloses job loss, missed payments and distress after repeated contact. An AI assistant drafts a standard response. Corentis Shield checks whether that output should pause, route to human review and record evidence before any customer communication proceeds.

The Corentis answer

Corentis Shield sits between AI agents and sensitive actions. AI proposes. Corentis checks. Lower-risk outputs can continue. Sensitive outputs pause, escalate or route to human review. Evidence is captured as the workflow runs.

Evidence that the problem is real

The evidence context in this pack shows large complaint volumes, material redress costs, vulnerable-customer pressure and fast AI adoption. Together, these signals point to a simple commercial truth: regulated organisations will need practical checkpoints before AI acts in sensitive workflows.

Why this can become infrastructure

The first wedge is narrow, but the reusable assets are broad: control schemas, scenario libraries, evidence artefacts, checkpoint logic and evaluation methods. The opportunity is runtime control infrastructure for regulated AI-agent workflows.

What makes Corentis investable

Corentis is building at the boundary where AI-generated intention becomes real-world action. That boundary is commercially valuable because regulated organisations need confidence, reviewability and evidence before they can scale AI agents in sensitive workflows.

What we are ready to prove

The next stage is validation. Corentis is ready to show how checkpointing can make AI-assisted workflows more reviewable, controllable and evidence-generating before sensitive customer actions move forward.

- Unsafe direct-action attempts caught.
- Vulnerable-customer escalation accuracy.
- Evidence completeness score.
- Human-review routing accuracy.
- False positive and false negative balance.
- Reviewer confidence and clarity of go/no-go decisions.

12-month proof plan

The next 12 months should turn the opportunity into evidence.

- V2 website and evidence pack live.
- ControlBench feasibility study prepared.
- Strategic R&D route prepared.
- Financial-services pilot route prepared.
- Benchmark scenario library defined.
- Baseline versus checkpoint workflow tests designed.
- Evidence completeness scoring prepared.
- Sample pilot reports produced.
- Design-partner conversations started.
- Route to controlled pilot clarified.

What strategic support would unlock

Strategic support would help Corentis deepen technical validation, build the scenario library, harden the evaluation framework, prepare design-partner pilots and turn the checkpoint concept into evidence-led market credibility.

Next conversation

If your organisation is exploring AI agents in regulated workflows, Corentis is ready for a focused conversation about validation, pilot design and strategic support.

SELECTED SIGNALS

Evidence context

FCA COMPLAINTS DATA

UK financial services firms received 1.85m complaints in 2025 H1.

Financial Conduct Authority, 23 October 2025

FCA COMPLAINTS REDRESS

Total redress in FCA complaints data was £283m in 2025 H1.

Financial Conduct Authority, 23 October 2025

FOS COMPLAINTS DATA

The Financial Ombudsman Service received 305,726 new complaints in 2024/25.

Financial Ombudsman Service, 2 July 2025

MCKINSEY GLOBAL AI SURVEY

88% of respondents in McKinsey's 2025 global survey reported regular AI use in at least one business function.

McKinsey & Company, 5 November 2025

MCKINSEY GLOBAL AI SURVEY

23% of respondents said their organisations are scaling an agentic AI system somewhere in the enterprise.

McKinsey & Company, 5 November 2025

IBM / PONEMON

63% of breached organisations lacked AI governance policies to manage AI or prevent shadow AI.

IBM / Ponemon Institute, 2025

Selected sources

Financial Conduct Authority: Aggregate complaints data: 2025 H1

Date/status: 23 October 2025. Source domain: fca.org.uk.
UK financial services complaints volume context.

Financial Conduct Authority: Aggregate complaints data: 2025 H1

Date/status: 23 October 2025. Source domain: fca.org.uk.
UK financial services complaints redress context.

Financial Ombudsman Service: Annual complaints data and insight 2024/25

Date/status: 2 July 2025. Source domain: financial-ombudsman.org.uk.
UK complaints escalation pressure context.

McKinsey & Company: The State of AI: Global Survey 2025

Date/status: 5 November 2025. Source domain: mckinsey.com.
Global cross-industry AI adoption context.

McKinsey & Company: The State of AI: Global Survey 2025

Date/status: 5 November 2025. Source domain: mckinsey.com.
Global agentic AI momentum context.

IBM / Ponemon Institute: Cost of a Data Breach Report 2025

Date/status: 2025. Source domain: ibm.com.
Security and AI governance-gap context.

Company details and next step

Corentis Shield is provided by Corentis Technologies Ltd. Company No. 17182737. Company type: Private limited company. Registered office: Suite A, 82 James Carter Road, Mildenhall, IP28 7DE, United Kingdom. Contact: hello@corentis.co.uk.

[Start a Conversation](#)