

Corentis Shield

AI checkpoint for regulated workflows

AI assurance stakeholders

Runtime Assurance Innovation Brief

Moving AI assurance closer to the action boundary.

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

A public assurance brief explaining how Corentis can create runtime evidence before sensitive AI-assisted actions move forward.

Generated April 2026

For discussion and pilot exploration only

Overview

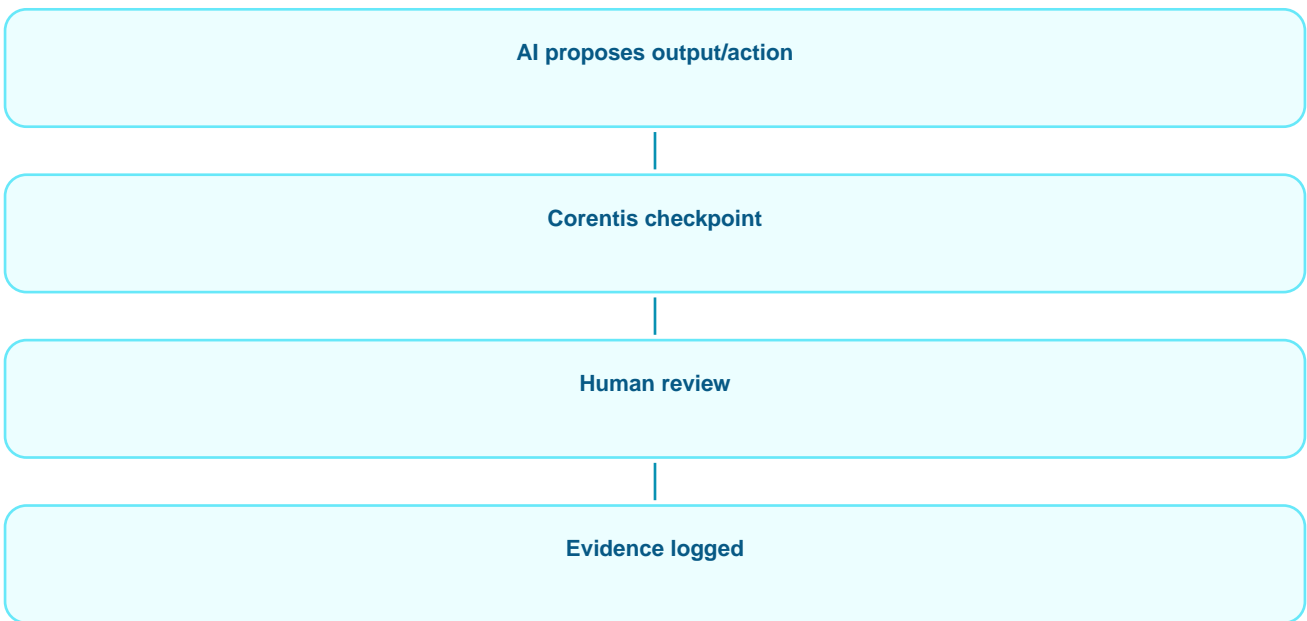
Core position

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.

Assurance should not only happen after the event

Many assurance approaches focus on policies, model evaluation or post-event review. Those are important, but they do not fully answer what should happen when AI is about to act in a sensitive workflow.

The action boundary

The action boundary is the moment when an AI-generated output could affect a customer, case file, decision or operational workflow. Corentis is designed to make that moment visible and controllable.

How Corentis creates runtime evidence

Corentis checks proposed AI outputs against policy intent, risk rules, human-review requirements and evidence requirements. The decision to proceed, pause, escalate or review becomes part of the evidence trail.

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.

What assurance reviewers could inspect

Reviewers could inspect the scenario tested, the proposed output, the triggered control, the checkpoint decision, the review route, the policy version and the evidence completeness score.

Evidence outputs

The evidence layer can include a control matrix, scenario test log, blocked-action log, review queue history, policy version trail, evidence completeness score, pilot report and go/no-go recommendation.

Validation metrics

A validation project should measure scenario coverage, policy-to-control mapping completeness, unsafe direct-action attempts caught, human-review routing accuracy, blocked-action explainability and audit artefact completeness.

Why this matters for trustworthy AI adoption

Trustworthy AI adoption in regulated workflows needs evidence at the point where action is about to happen. Corentis is building for that moment.

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

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

MCKINSEY GLOBAL AI SURVEY

51% of respondents from organisations using AI said their organisations had seen at least one negative consequence.

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

SALESFORCE AI CUSTOMER RESEARCH

72% of customers say it is important to know if they are communicating with an AI agent.

Salesforce, 2026 page accessed / report current at access

Selected sources

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.

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

Date/status: 5 November 2025. Source domain: mckinsey.com.
Global AI risk and consequence context.

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

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

Salesforce: State of the AI Connected Customer

Date/status: 2026 page accessed / report current at access. Source domain: salesforce.com.

Vendor AI trust and customer expectations 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](#)