AI Assurance Needs Operational Evidence

You can test an AI model, define governance policies and establish controls. But how do you know what actually happens when AI agents operate inside your business?

Evaluation tests what AI can do; governance defines what should apply. Lusinova shows what actually happens in operation. **Unlike AI observability alone, Lusinova can connect available AI-agent execution evidence with enterprise processes and human activities, revealing how agent behaviour affects actual business operations.**

AI agent execution + Enterprise-system events + Human activities → Evidence of actual AI operations

01 WHAT LUSINOVA CAN ASSESS

AI Agent Behaviour & Traceability
Reconstruct agent execution, tool interactions and operational handoffs across systems.

Operational Risk & Deviation Identification
Identify unexpected behaviour, failures, repeated loops and process deviations.

Human Oversight & Control Adherence
Assess evidence that required reviews, approvals and escalations took place.

Incident Investigation & Root-Cause Analysis
Investigate failures, identify potential contributing factors and assess their operational consequences using available execution evidence.

Continuous Operational Monitoring
Track execution patterns, exceptions and changes in behaviour over time.

Operational Impact & Improvement Verification
Link agent behaviour to outcomes and benchmark the effect of changes.

02 SUPPORTING AI GOVERNANCE, RISK AND ASSURANCE

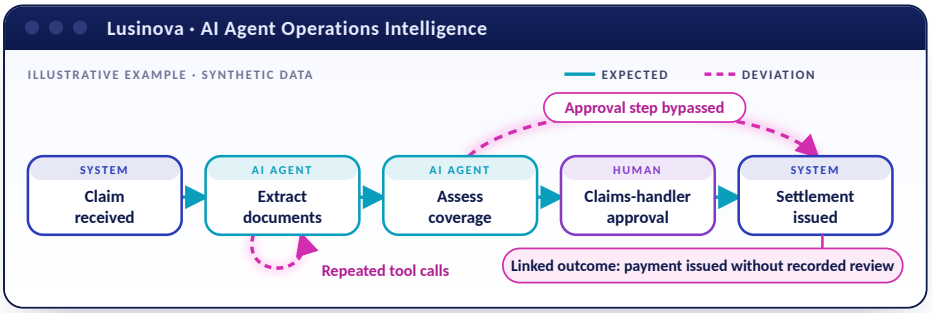
Lusinova provides operational evidence to help organisations:

- Investigate whether AI agents follow approved operating procedures and controls.
- Identify deviations, failures and unexpected behaviour.
- Examine human oversight, approvals and escalation.
- Support audits, incident investigations and ongoing operational monitoring.

Lusinova complements model-level safety evaluations by providing evidence of AI behaviour, operational controls and human oversight in real business environments.

Lusinova does not replace assurance providers or compliance assessments; it supplies the operational evidence that helps make AI assurance possible.

03 FROM AI AGENT EXECUTION TO ASSURANCE EVIDENCE



1 Reconstruct
Rebuild the actual sequence across agent, system and human steps.

2 Identify deviations
Flag departures from the defined process, such as a skipped approval or a repeated loop.

3 Connect to outcomes
Link each deviation to its operational consequence, such as a payment issued without review.

30-Day AI Operational Assurance Discovery

Understand how your AI agents actually behave — and where operational risks may exist. A focused assessment of selected AI agents and their interaction with business operations.

Proposed scope Up to two AI agents, one enterprise workflow and agreed operational evidence. Confirmed at scoping; needs approved access to traces and system records. Major integration excluded.

Examples of what we investigate:

- AI agents bypassing required approval steps.
- Unexpected tool calls or system interactions.
- Repeated failures, execution loops and uncontrolled escalations.
- Missing evidence of required human intervention.

Illustrative scenarios, not automatic detection of every such risk.

In 30 days, subject to data readiness:

1. Reconstruct actual agent behaviour and execution paths.
2. Identify deviations, failures and potential control gaps.
3. Assess available evidence of human oversight and escalation.
4. Investigate selected risks and operational consequences.
5. Deliver an assurance findings report and prioritised recommendations.

WHAT YOU RECEIVE

Execution Evidence How AI agents actually operated.

Risk & Control Findings Deviations, exceptions and evidence gaps.

Operational Assurance Evidence Report Findings on oversight, controls and incidents.

Action Plan Prioritised recommendations and verification criteria.

Discuss AI Operational Assurance with Lusinova

A 30-minute demonstration and discussion of enterprise applications.

a.gevorkian@lusinova.com
lusinova.com