

CODANCE PUBLICATION

Engineering Case, Hypothesis and Alpha Statement

The intellectual foundation, testable proposition and bounded status of Project Aurora

Control	Value
Document	COD-PUB-THS-001 v0.1 25 August 2026
Owner	Codance Ltd
Platform	Project Aurora working alpha
Purpose	Public technical position paper
Status	First substantive draft for competent review
Review discipline	Distinguish demonstrated evidence, engineering proposition and programme intent

CENTRAL PROPOSITION

Agentic AI can be appropriately assured only through an evidential process: material actions, decisions, delegations and consequences must remain attributable, reconstructible and open to competent challenge, while the assurance apparatus is subjected to equivalent evidential discipline.

Purpose and reading discipline

This paper sets out Codance's engineering case for how agentic AI can be appropriately assured, structured as three things that are easily blurred: a thesis — Codance's argued position on the assurance problem; a testable hypothesis about a practical response; and the narrower evidence presently supplied by the Aurora working alpha. The thesis is an argued position. The hypothesis is a proposition that can be challenged experimentally. The alpha is an implemented instrument whose demonstrated behaviour is evidence for some parts of that proposition, not proof of the whole

CLAIM DISCIPLINE

The existence of a working platform does not establish that every agentic system is observable, that every assurance method is valid, or that Aurora itself is an accredited assurance instrument. Each conclusion remains bounded by the commissioned system, method, evidence and authority.

Layer	Question	Present status
Thesis	What conditions must valid assurance satisfy?	Argued engineering position
Hypothesis	Can an adaptive campaign meet those conditions in operation?	Testable proposition
Alpha	Has Codance implemented and exercised relevant mechanisms?	Demonstrated within declared boundaries

1. The engineering thesis

Agentic AI does not operate as an isolated model. It acts through applications, tools, policies, data, people and infrastructure. A plausible final output cannot, by itself, establish that the path to that output was authorised, complete or safe. The material object of assurance is therefore the operating system formed by the probabilistic agent and the human and technical arrangements through which it acquires consequence.

Appropriate assurance requires more than observation. Evidence must be attributable to defined identities and boundaries; sufficiently complete for the stated question; preserved in a form that supports reconstruction; and accessible to a person or institution with legitimate authority to interpret it. A method that reaches a favourable answer while concealing its own influence is not evidentially adequate.

Every material action, decision, delegation and consequence of the agent and containing application must be attributable.

- The declared System Under Test, configuration, assurance method, collection boundary and execution identity must be controlled.
- The assurance procedure's own prompts, transformations, tools, model calls, tests and omissions must be reconstructible.
- Missing evidence must remain missing; it must not be silently replaced by inference or expected behaviour.
- A competent human authority must retain responsibility for the meaning and use of the finding.

THESIS

Agentic AI can be appropriately assured only through an evidential process. Declared boundaries, equivalent evidential rigour for the assurance apparatus and legitimate human authority are prerequisites to a valid assurance outcome.

1a. Why ordinary output checking is insufficient

Output checking answers whether a visible result appears acceptable. Assurance must also address how that result came about, which constraints applied, what information influenced it, which tools were invoked, whether consequential requests occurred, and whether the apparent outcome masks a materially different process. For example; two executions display the same final refusal while one attempted an unauthorised payment. Conversely, two executions receive equivalent source records yet reach different recommendations because a commissioned evidence guide changed what was presented to a model.

Visible question	Assurance question
Was the final answer reasonable?	Was the answer produced through the commissioned and authorised pathway?
Did the “something” complete?	Was any “something” request attempted before authority was established?
Did the model receive the intended record?	Did configuration alter the evidence selected, framed or supplied to the model?
Did the workflow finish?	Can its material events and the assurance procedure be reconstructed and challenged?

1b. Conditions of valid assurance

Codance's thesis treats assurance validity as conditional rather than absolute. A finding is valid only for the question, system, method, execution and evidence for which it was made. This rejects both universal claims from bounded tests and the idea that volume of telemetry can compensate for an undefined assurance question.

Condition	Operational meaning	Failure consequence
Declared boundary	The controlled article, interfaces, configuration and exclusions are explicit.	The finding cannot be reliably scoped.
Evidential sufficiency	Required observations are available with necessary identity, timing and integrity.	The criterion becomes UNKNOWN or the test is not commissioned.
Equivalent rigour	Material influences of the assurance apparatus are retained and attributable.	The procedure cannot support independent challenge.
Legitimate authority	A competent authority owns interpretation and consequential use.	The platform would be substituting machinery for accountable judgement.
Challenge-ability	Evidence, transformations, limitations and deviations can be examined.	The conclusion is an assertion rather than an assurance finding.

2. The testable hypothesis

HYPOTHESIS

An adaptive assurance campaign can combine deterministic and probabilistic methods while preserving sufficient external evidence for a competent reviewer to reconstruct and challenge both the material operation of the System Under Test and every material influence exerted by the assurance procedure itself.

The hypothesis is intentionally falsifiable. A Codance research project would not support merely producing reports or retaining logs. The evidence must show that controlled identity survives routing and processing; deterministic controls remain exact; probabilistic contributions are bounded and attributable; completed history remains immutable; adaptive changes affect only authorised future work; and the report can be traced backwards to registered source evidence.

Support would require	Contrary evidence would include
Independent reconstruction of material SUT and method influences.	A material action or method influence that cannot be attributed.
Controlled comparison separates declared variables from invariants.	Uncontrolled differences presented as causal evidence.
Missing or malformed evidence produces refusal, failure or UNKNOWN.	Silent repair, enrichment or assumption of required evidence.
Human authority remains explicit at interpretation and action boundaries.	An automated conclusion represented as legitimate authority.
The apparatus exposes its own versioned configuration and lineage.	A favourable result whose procedure cannot be independently examined.

3. What the working alpha demonstrates

Codance’s project Aurora, as an alpha product, implements a campaign-oriented evidence architecture across distinct operational and evidential components. Governed transmissions are admitted through registered identities; inputs are processed through versioned routes, stages and skills; work-in-progress records preserve dependencies and products; observations are associated with collection points and execution identities; and reports retain lineage to the evidence from which they were derived.

On 24 August 2026 the ClearLedger controlled pair demonstration exercised this pathway end to end. Two executions used the same sealed SUT (system under test) revision, fixture, archived source records and pinned model identity. The declared difference was the “commissioned receipt-evidence guide”. Aurora retained 18 operational-journal signals for each execution, generated individual execution assessments, accepted a human selection of the pair, gathered the registered lineages, tested the invariants, identified the first material divergence and emitted a human-readable pair report.

- Control configuration: ERP Goods In authoritative; recommendation RECOMMEND_HOLD.
- Comparison configuration: operational correspondence; recommendation RECOMMEND_PAY.
- First material divergence: the model-request material after common archive acquisition.
- Pair conclusion: CONTROLLED_CONFIGURATION_EFFECT_SUPPORTED.
- Safety boundary: human authority retained; no payment attempt or execution recorded.

WHAT THIS SUPPORTS

The demonstration supports the narrower claim that Aurora can execute a declared controlled comparison, retain both system and procedure evidence, and produce a reconstructible attributed report for that commissioned case.

3a. What has not yet been established

- Generality across arbitrary agentic systems, assurance methods or sectors.
- Completeness of observation where a SUT does not expose the required surfaces.
- Accreditation, metrological calibration or conformity with a particular external standard.
- Predictive validity of any probabilistic assessment beyond the demonstrated fixture.
- Operational scale, resilience, security and supportability required for commercial deployment.
- That evidence capture alone determines the meaning, acceptability or consequence of a finding.

These are programme questions, not footnotes. Codance’s next phase is designed to submit the method and instrument to competent external challenge, operate bounded real-world beta cases, and convert demonstrated mechanisms into qualified and supportable infrastructure.

PRESENT ALPHA STATEMENT

Project Aurora is a working alpha platform for modelling and operating bounded assurance campaigns as attributable evidence-producing processes. It is suitable for controlled technical examination and representative beta development; it is not presented as an accredited assurance instrument or universal solution.

4. Review questions

- 1. Are the thesis conditions necessary, sufficient and expressed at the correct level?
- 2. Is the hypothesis genuinely falsifiable, and what experimental evidence would be decisive?
- 3. Which classes of assurance question cannot be supported through external operational evidence?
- 4. What calibration, qualification or reference artefacts would be required before Aurora could be treated as an assurance instrument?
- 5. Which parts of the ClearLedger result are demonstrative of platform capability, and which remain case-specific?

Codance invites challenge to these questions because the purpose of the programme is not to defend an attractive narrative. It is to determine, through controlled engineering work, where the thesis is useful, where the hypothesis holds, and where the instrument or method must change.

Document control and review

This is a first substantive public draft prepared for structured review. Before release, Codance will verify technical identities and claims against the controlled evidence, resolve reviewer comments, assign the approved publication version, regenerate the PDF and record the final file SHA-256.

Review question	Required disposition
Is every statement clearly classed as demonstrated fact, engineering proposition or programme intent?	Confirm, narrow, source or remove.
Can a competent reader identify the controlled boundary and prohibited over-claims?	Make the boundary explicit.
Do technical identities and results match retained evidence?	Verify against controlled records.
Does the document disclose material limitations and credible contrary interpretations?	Add or strengthen limitations.
Is the document useful without access to confidential implementation detail?	Improve explanation or supporting public evidence.