CODANCE PUBLICATION

From Capable AI to Justified Confidence

The assurance problem Codance addresses and the programme built to examine it

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

THE GAP

Capability tells us what an AI-enabled system can appear to do. Justified confidence requires evidence of what it did, under which authority, through which path, with what consequence and with what limitations.

Executive argument

Artificial intelligence is moving from producing content to participating in operational systems. As agents acquire tools, memory, delegation and access to consequential processes, the important question changes. It is no longer enough to ask whether a model can produce a convincing answer. Organisations must determine whether a particular system operated within an authorised boundary, whether its material influences and actions can be reconstructed, and whether a competent authority has adequate evidence for the confidence being claimed.

Codance's proposition is that this bridge from capability to confidence is an engineering and evidence problem. Project Aurora is the working alpha built to test that proposition. It models assurance methods as controlled campaigns, operates deterministic and probabilistic work through attributable processing, gathers evidence at declared collection points and produces reports whose conclusions remain linked to their source material and limitations.

NOT A CERTIFICATION CLAIM

Aurora does not turn the uncertain behaviour of AI into certainty. It is designed to make the basis, boundaries and residual uncertainty of a finding explicit and open to competent challenge.

1. Capability is not accountability

Modern models can summarise, plan, classify, call tools and generate software with impressive fluency. Those capabilities are valuable, but capability demonstrations usually optimise for task success. Operational accountability requires a different body of evidence: who initiated the work; what system and configuration were active; what information entered the process; which transformations and tools influenced it; what material requests were attempted; what actually happened; and who had authority to accept the result.

A human reviewer cannot responsibly answer those questions from a final chat transcript or application status alone. Nor is an instruction to keep a human 'in the loop' sufficient if the human receives only a polished conclusion. Human authority becomes meaningful when the reviewer can inspect evidence that is proportionate to the consequence and can distinguish observation from inference, deterministic fact from probabilistic interpretation, and declared absence from missing data.

Capability evidence	Confidence evidence
Benchmark score or successful task completion	Attributable operation within a declared boundary
Plausible visible output	Reconstructible material pathway and consequences
Model or application documentation	Controlled runtime identities, versions and evidence lineage
Human approval button	Competent authority supplied with adequate evidence that is open to challenge
Large telemetry volume	Evidence selected against a stated assurance question

2. The system is larger than the model

Consequential behaviour arises from an assembly: model, prompt and guidance; application logic; tools and permissions; retrieved data; policy and deterministic controls; orchestration; operators; external services; and the environment in which actions take effect. Assurance restricted to the model can miss defects introduced by any of the other components, including the assurance method itself.

- A correct policy may be evaluated after an irreversible request has already been emitted.
- Equivalent records may be framed differently by a configuration guide and drive different recommendations.
- A safe final refusal may conceal an earlier unauthorised attempt.
- A well-performing model may be embedded in an application whose identities, retries or tool boundaries are not governed.
- An assurance procedure may manufacture confidence by selecting convenient evidence or omitting its own influence.

The assurance object must therefore be explicitly bounded and must include the operational interfaces through which probabilistic behaviour becomes consequential. This does not require inspecting every internal model activation. It requires sufficient external evidence at the points where actions, decisions, delegations and consequences acquire meaning.

3. Assurance as an evidential process

Codance treats assurance as a process that produces a bounded finding from declared evidence under legitimate authority. The method may combine deterministic tests, probabilistic analysis, expert interpretation and human judgement, but the role of each must remain distinguishable. The process is valid only if a reviewer can understand what was observed, what was computed, what was inferred, what was unavailable and how those elements support the stated conclusion.

1. Define the assurance question and prohibited over-claims.
2. Control the SUT, method, execution, configuration and observation boundaries.
3. Admit only governed transmissions and preserve failures and refusals.
4. Execute versioned steps whose dependencies and products remain explicit.
5. Retain source evidence and material influences of the assurance procedure.
6. Apply deterministic and probabilistic methods according to declared roles.
7. Issue a finding with lineage, limitations, authority and a route for challenge.

EVIDENCE BEFORE INTERPRETATION

Aurora is intended to preserve the material from which measurement, assessment or judgement may be derived. It does not claim that evidence capture is itself measurement or that the platform owns the competent meaning of the result.

4. Project Aurora

Aurora is Codance's implementation of this operating idea. It separates campaign identity and admission, campaign-blind processing, device-bound skills, evidence retention and report construction. A governed delivery enters through one input plane. Its registered transmitter, transmission definition and address establish the permitted meaning and route. Core creates a canonical message identity, then executes the published route as a series of work records. Each step consumes declared dependencies, emits typed products and completes or fails with an attributable disposition.

External sensors can observe declared SUT collection points without giving the SUT authority over the assurance record. Campaign Register provides independent attribution. The evidence store retains products by identity and hash. Reports are generated as outputs of the same governed process and preserve lineage back to the relevant inputs and retained products.

Aurora contribution	Confidence question supported
Governed admission and attribution	Was this delivery authorised for this campaign and collection point?
Versioned routes, stages and skills	Which procedure actually operated?
Work and product lineage	Which prior evidence and transformations influenced this output?
Sensors and collection points	What externally observable event or consequence was retained?
Typed terminal outputs and failures	Did processing finish, refuse or fail, and how?
Human-readable report with hashes	Can the finding be traced, reviewed and challenged?

5. Demonstrated controlled difference

The ClearLedger demonstrator asked whether changing only a commissioned receipt-evidence configuration could alter a recommendation produced from otherwise equivalent records and model identity. The control treated ERP Goods In as authoritative and recommended HOLD. The comparison treated operational correspondence as relevant and recommended PAY. Aurora verified the declared invariants, isolated the configuration difference and identified the first material divergence in the model-request material after common archive acquisition.

The demonstration matters because it exposes a class of assurance problem that a final-output view can miss: the visible recommendation is influenced not only by the source records and model, but by the evidence-selection and framing procedure. Aurora retained both the SUT execution evidence and the assurance procedure's contribution, allowing the difference to be attributed within the bounded test.

Controlled invariant	Controlled difference	Observed consequence
Invoice, sealed SUT revision, archived records and value hashes, pinned model, campaign contract, human authority	ERP-authoritative versus operational-correspondence evidence guide	RECOMMEND_HOLD versus RECOMMEND_PAY

DEMONSTRATED CONCLUSION

CONTROLLED_CONFIGURATION_EFFECT_SUPPORTED. This is a commissioning-demonstrator result, not a certification of ClearLedger or proof of general causal validity.

6. From alpha to justified use

A credible assurance platform must earn confidence in stages. Codance's working alpha establishes an implementable foundation. The next work is to expose the thesis, method, apparatus and reports to competent external challenge; qualify the instrument against known reference cases and failure modes; operate bounded beta systems with real authorities and consequences; and convert the resulting learning into supportable, secure and repeatable infrastructure.

- Independent examination of terminology, evidence sufficiency, reconstruction and claim discipline.
- Reference cases that test success, failure, refusal, missing evidence, collision, retry and method deviation.
- Bounded beta engagements with declared SUT surfaces and legitimate finding authorities.
- Security, resilience, operational support and controlled release engineering.
- Assurance-method development that remains evidence-gated rather than outcome-seeking.

The objective is not to declare trust in AI. It is to make particular confidence claims better justified, more precisely bounded and more open to challenge.

7. Invitation

Codance invites three forms of participation: competent technical and assurance challenge; bounded real-world systems on which meaningful questions can be tested; and aligned investment in a defined programme of product development, independent examination and beta operation.

The strongest collaborators will bring more than one kind of value: measurement or assurance expertise, system access, domain authority, security or standards knowledge, market access, and the willingness to test whether the evidence really supports the claim. The appropriate first step is a short technical conversation about the question, the system and the evidence that would matter.

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.