

CODANCE PUBLICATION

Aurora Development Pathway

From demonstrated working alpha through independent challenge and bounded real-world beta

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

PROGRAMME PRINCIPLE

Aurora development is evidence-gated rather than calendar-gated. Each phase must produce defined technical and evidential outputs before the next claim, deployment boundary or commercial commitment is authorised.

Purpose and programme logic

This paper describes how Codance intends to progress Aurora from a founder-operated working alpha into a qualified, supportable platform for bounded assurance and operational-evidence engagements. It is a development pathway, not a promise that maturity follows automatically from elapsed time or investment.

The programme combines technical development, independent challenge and real-system learning. Product engineering without external examination risks polishing an untested theory. Examination without representative systems risks validating only laboratory behaviour. Beta operation without controlled infrastructure risks confusing delivery pressure with evidence of readiness.

GATE RULE

No phase is complete because its activities occurred. It is complete only when its evidence package supports the defined exit decision and unresolved limitations have an authorised disposition.

1. Established working-alpha baseline

The present baseline includes governed input admission and campaign attribution; versioned routes, stages and device-bound skills; DAG work and product lineage; external operational-journal sensors; deterministic and model-assisted assessment; individual execution closure; human selection of a completed control/comparison pair; comparative assessment; and human-readable reporting.

The ClearLedger demonstration exercised this integrated pathway on one sealed SUT revision with two commissioned configurations. All reporting steps completed and the pair report retained the controlled invariants, difference, first material divergence, conclusion, limitations and product hashes.

Baseline capability	Evidence presently available
Governed admission	Campaign Register admission, submission and binding receipts
Attributable processing	Published route and completed DAG-step records
Independent observation	S-03 sensor checkpoints, config locks and retained signals
Execution reporting	Completed evidence snapshots, assessments and reports
Controlled comparison	Human selection, pair snapshot, assessment and report
Claim discipline	Explicit commissioning-demonstrator limitation

BASELINE LIMIT

The alpha demonstrates integrated mechanisms in a controlled case. It is not yet qualified for arbitrary SUTs, production service levels or accredited assurance use.

2. Phase A — independent technical challenge

The first development phase submits the thesis, commissioning method, evidence model, instrument behaviour and ClearLedger records to people with relevant competence and no obligation to preserve Codance's preferred interpretation. The objective is to identify invalid assumptions, missing evidence, ambiguous terminology and inadequate qualification requirements.

- Review the distinction between evidence capture, measurement, assessment and finding authority.
- Challenge the sufficiency and independence of Campaign Register and sensor attribution.
- Attempt reconstruction from report to source without relying on developer explanation.
- Review failure semantics, evidence absence, method deviation and UNKNOWN outcomes.
- Define candidate reference cases, calibration concepts and qualification evidence.
- Assess security and adversarial assumptions for a potentially hostile SUT.

Exit evidence	Gate question
Recorded review findings and dispositions	Are material criticisms visible and controlled?
Revised thesis, hypothesis and method boundary	Is the proposition still coherent and falsifiable?
Qualification and reference-case plan	Can the instrument be tested against known conditions?
Prioritised architecture defect ledger	Are changes driven by evidence rather than preference?
External challenge report or documented review	Is there competent evidence beyond Codance's own assertion?

3. Phase B — qualified instrument and reference suite

Before Aurora is trusted to examine consequential systems, its own apparatus must be challenged under known conditions. Codance will build a reference suite that covers successful operation and designed failure: malformed and missing evidence, duplicate delivery, retry, identity collision, gate closure, route mismatch, sensor discontinuity, skill exception, failed dependency, incomplete terminality and report-lineage break.

Qualification is configuration-specific. A sensor, adapter, route or skill does not become universally qualified because one version passed one test. Evidence must bind the result to the tested instance, version, environment, configuration and acceptance criteria.

Qualification family	Representative question
Identity and admission	Does unauthorised or mismatched input fail closed and remain recorded?
Integrity and order	Are hash, sequence, duplicate and discontinuity defects detected?
Process and terminality	Does every admitted input reach an attributable final disposition?
Method behaviour	Do known evidence conditions produce the specified PASS, FAIL, UNKNOWN or deviation state?
Reconstruction	Can an independent reviewer reproduce the path and product identities?
Operational resilience	Are restart, retry, outage and partial-service behaviours controlled?

4. Phase C — bounded real-world beta

Beta engagements apply Aurora to representative systems whose owners can define a meaningful assurance or operational-evidence question and expose sufficient surfaces lawfully. The beta is not an enterprise-wide promise to assure AI. Each case is a bounded Campaign with explicit exclusions, legitimate authority, controlled change and an agreed evidence-access model.

Admission criterion	Why it matters
Defined system and version	Prevents findings drifting across uncontrolled change.
Meaningful consequential question	Ensures evidence serves a real decision rather than a technology demonstration.

Admission criterion	Why it matters
Observable collection points	Makes the required method evidence technically obtainable.
Competent finding authority	Keeps interpretation and use with a legitimate human or institution.
Safe stimulus and rollback	Controls risk introduced by the assurance activity.
Evidence rights and retention	Allows reconstruction, challenge and lawful handling.
Willingness to expose failure	Prevents the beta becoming a favourable-outcome exercise.

Beta outputs include a commissioned method, SUT–Method map, installed and qualified apparatus, execution evidence, report, limitation register, client review and reusable implementation learning.

5. Phase D — deployable infrastructure

The alpha estate proves concepts but commercial operation requires engineering for repeatability and support. This includes tenant and campaign isolation, secure deployment, identity lifecycle, secrets, backup and restoration, availability, monitoring, controlled upgrades, evidence export, retention policy, incident response, installation tooling and documented operator roles.

- Reproducible deployment and environment qualification.
- Stable APIs and versioned schemas with migration discipline.
- Least-privilege access and auditable administrative operations.
- Supportable sensor, adapter and skill lifecycle.
- Client-controlled evidence access and sensitive-data boundaries.
- Operational service definitions, recovery objectives and incident evidence.
- Release documentation, training and competent handover.

PRODUCT READY TEST

A platform is not commercially repeatable merely because its developers can operate it. A competent delivery team must be able to commission, run, diagnose and support a bounded engagement from controlled documentation and tools.

6. Phase E — advanced assurance methods

Adaptive methods are a later capability, not a shortcut around foundational controls. Aurora may authorise a completed step to add or re-prioritise future work, select an appropriate analysis model or request additional evidence. Such adaptation is acceptable only when the authority, trigger, mutation, new dependencies and resulting products remain attributable and completed history is unchanged.

- Comparative and counterfactual campaign patterns.
- Evidence-led adaptive collection within pre-authorised boundaries.
- Multiple-model selection with retained decision rationale and identity.
- Method ensembles that separate deterministic, probabilistic and expert contributions.
- Longitudinal operational-evidence campaigns across controlled system change.
- Reference libraries for known failure patterns and reconstruction tests.

Each advanced method requires its own validity argument, qualification evidence and claim boundary. Platform flexibility does not validate the method executed on it.

7. Twelve-month delivery frame

Period	Primary work	Decision evidence
Months 1–3	External challenge, defect closure, reference-suite design	Reviewed method boundary and qualification plan
Months 3–6	Instrument qualification, security and deployment hardening	Reference results and controlled release baseline
Months 5–10	Two or more bounded beta cases	Commissioned reports, client challenge and reusable patterns
Months 8–12	Support model, evidence access, product packaging and next-method work	Deployable beta capability and next-stage investment decision

Work streams overlap where evidence permits. The schedule is deliberately subordinate to gate quality. A failed gate produces remediation or a revised claim—not a relabelled milestone.

8. Programme risks and controls

Risk	Control
Over-claiming from one demonstrator	Publication review, explicit claim boundaries and external challenge
Instrument influences the SUT	Declared stimuli, independent collection points and change control
Insufficient SUT observability	Commissioning compatibility map and refusal to manufacture evidence
Founder dependency	Documented methods, installation tooling and team-deliverable operations
Sensitive evidence exposure	Client-side retention patterns, minimisation and governed access
Commercial pressure distorts findings	Pre-agreed method, immutable execution history and legitimate authority
Complexity outruns supportability	Evidence-gated feature admission and bounded beta scope

9. Participation and success

Codance is inviting aligned investment to fund a defined programme rather than an open-ended research allowance. Capital is necessary but not sufficient in itself. The programme also requires competent technical challenge, representative systems, domain authorities, security and standards expertise, and partners able to help convert a founder-operated alpha into a repeatable specialist product.

Success after twelve months would not be a universal claim that Aurora assures AI. It would be a stronger and independently challenged thesis, a qualified instrument baseline, evidence from bounded real-world cases, a supportable deployment model and a defensible decision about the platform's next commercial and technical scope.

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.