PUBLIC EXECUTIVE SUMMARY · 03 IN READING ORDER

OUTLIER

/ AI · REDESIGN EXECUTIVE SUMMARY

PLATFORM REPLACEMENT LOGIC

The corrected primitive, minimum replacement, and proof standard.

> *"Outlier is recruiting experts to do compliance work.*
>
> *That is why nothing else has worked."*

ARTIFACT FAMILY

#	ARTIFACT	ACCESS
01	Translation Artifact	Public
02	Diagnostic Artifact	Public
03	**Redesign Executive Summary**	← *This document*
04	Redesign Artifact	Restricted

PREPARED BY Jamie Forrester · Independent Systems Architect · Edinburgh, UK
CONTACT hello@jamieforrester.com · jamieforrester.com
DATE April 2026
CLASSIFICATION Public · 03 in reading order · not for implementation

SECTION 0

WHAT THIS DOCUMENT IS

This is the public executive summary of a structural replacement specification. It is not a roadmap, a reform proposal, an advisory document, or an optimisation programme. It summarises the corrected primitive, the non-negotiable continuity object the redesign must preserve, the target-state architecture, what must stop, and what the board is approving.

The governing diagnosis is established in the Diagnostic Artifact. This document does not reopen it. Every design decision in the full Redesign Artifact is traceable to that diagnosis.

What this document contains: corrected primitive, truth kernel, target-state architecture, minimum replacement conditions, dependency order, fake-progress test, consequence of deferral, board verdict.

What the full Redesign Artifact adds: the complete six-move transition sequence; rules and refusal invariants including the Compensation Truthfulness Invariant; drift immune system with fail-closed primitive regression gates; nine admissibility conditions with sign-off requirements; pilot validity rules and failure criteria; Consequence Review Trigger Matrix; Queue-State Visibility Model; Expert Signal Lifecycle; object-level builder handoff; automated testability requirements; and Appendix B binding Redesign Acceptance Protocol. The full Redesign Artifact is restricted and not published publicly.

Important: the full Redesign Artifact is released only after the client signs Appendix B — the binding Redesign Acceptance Protocol. That protocol commits the organisation to the dependency order, the Vanish List, and the minimum replacement conditions before any build begins.

> ARTIFACT FAMILY — READING ORDER
>
> ***Recommended reading order: Translation Artifact (best first read) → Diagnostic Artifact (full evidence) → This document (replacement logic) → Full Redesign Artifact (restricted; not published). Public artifacts demonstrate the diagnostic and replacement logic. The restricted artifact specifies the implementation.***

Judgment boundary

This document makes no claim about intent, bad faith, undisclosed internal systems, or adjudicated legal liability. It summarises a structural redesign derived from the public Diagnostic Artifact, which is itself based entirely on publicly available evidence.

What this document does not say

- → It does not say algorithmic scoring must be removed — it says scoring must not directly trigger irreversible consequence without human review
- → It does not say Queue Managers should be removed — it says their role must change from load-bearing coherence to supplementary escalation
- → It does not say all work should become expert judgment work
- → It does not say Outlier must guarantee continuous work availability
- → It does not say the contractor classification model must change

SECTION 0.1

REDESIGN SPINE

The complete replacement specification stated in governed form. Every section of this document and the full Redesign Artifact traces back to one or more rows of this table.

REDESIGN ELEMENT	SPECIFICATION
Governing mismatch	Single contributor frame coordinating two incompatible work classes under one gate, one scoring model, and one consequence path
Corrected primitive	Dual-stream contributor coordination system
Truth kernel	Portable cross-project contributor continuity record — the one object the redesign must not sacrifice
Stream A	Constrained reliability execution — rubric-conformant, variance-suppressed, high-throughput
Stream B	Bounded expert judgment — specialist reasoning, correct variance valued, upstream reframing authority
Stream contracts	Each stream has its own promise, measurement, and compensation. Reliability work cannot be advertised as expert judgment. Expert judgment cannot be paid as throughput labour.
First move	Primitive declaration at leadership level — all downstream design depends on this
Minimum replacement	Separate gates + continuity record + queue visibility + consequence review + expert signal channel. All five. If any are missing, the change is not a redesign. (See Section 5 for details.)
Proof condition	Same pressure → different response
Primary recapture risk	Expert stream becomes symbolic — exists in language without authority to change task definitions

SECTION 1

GOVERNING DIAGNOSIS

GOVERNING DIAGNOSIS — LOCKED

Outlier recruits through the grammar of an expertise market while operating through the mechanics of a reliability pipeline. In plain terms: the platform is recruiting experts to do compliance work. That is why nothing else has worked. The platform's architecture derives legitimacy benefits from expert-coded recruitment while operationally consuming compliance-governed output. Interface-level corrections cannot resolve this. The failure originates at the primitive.

Every symptom — expert contributors failing production, quality scores not predicting output quality, Queue Managers becoming permanent infrastructure, community tooling filling information gaps the platform should carry — follows directly from this one structural condition.

The correction is not to make the single frame clearer, stricter, or better supported. The correction is to stop using one frame for two incompatible things.

SECTION 2

TRUTH KERNEL

The non-negotiable continuity object this redesign must preserve is a portable contributor continuity record — a persistent, cross-project object that preserves demonstrated performance, qualification history, contribution state, consequence history, dispute history, and routing-relevant context.

The governing rule: demonstrated contributor value must travel with the contributor unless there is explicit, rule-based cause for it not to. The system must be able to learn from its own best contributors over time.

The non-negotiable test: if contributor continuity remains fragmented and per-project at the end of the redesign, the redesign has failed regardless of any surface improvement.

Visibility requirement: the continuity record must be visible enough to be contested. It cannot become another hidden file. A contributor must be able to see what the system knows about them and request formal review of any entry.

SECTION 3

CORRECTED PRIMITIVE

Malformed primitive: one contributor market purchasing "expert judgment" while actually consuming variance-controlled instruction adherence at scale.

Corrected primitive: a dual-stream contributor coordination system. One stream coordinates constrained reliability execution. The other coordinates bounded expert judgment. Each stream has its own intake signal, qualification gate, routing logic, evaluation logic, and consequence path. The two streams are coordinated but not conflated.

Stream contracts

Each stream must have its own explicit contract: what the platform promises, what it measures, and how it compensates.

DIMENSION	STREAM A — RELIABILITY EXECUTION	STREAM B — EXPERT JUDGMENT
Promise to contributor	Constrained reliability work: bounded tasks, rubric adherence, consistent metrics	Bounded expert judgment: specialist reasoning, rubric challenge, upstream influence on task design
What is measured	Rubric conformance, throughput, consistency under constraint	Signal quality, reasoning depth, rubric challenge quality, upstream impact
Compensation logic	Per-task or per-time aligned to execution; paid onboarding where required	Per expert contribution; reflects specialist input value, not throughput volume. Compensation convergence toward Stream A rates is the primary compensation recapture signal (Fake Progress Test, card 6).
What variance means	Suppressed — variance is the failure signal	Valued — correct variance is the product

Contract rule: *reliability work cannot be advertised as expert judgment. Expert judgment cannot be paid and measured as throughput labour. A contributor must be recruited under and evaluated against the same contract.*

SECTION 4

TARGET-STATE ARCHITECTURE

The dual-stream design separates what the current system has collapsed. The shared infrastructure connects both streams without conflating them. In this document, irreversible consequence means any action that materially affects contributor access, pay, tier status, project eligibility, routing priority, or account standing and cannot be automatically reversed by the contributor without formal review.

Function allocation

FUNCTION	CURRENT STATE	TARGET STATE
Sensing	Algorithmic scoring calibrated to rubric adherence proxies, not actual work quality	Stream-specific scoring calibrated to actual stream quality signals
Interpretation	Fragmented across informal QM layer and shifting guideline documents	Bounded formal interpretation layer with design authority
Authority	Fused with sensing — score-mediated signals appear capable of triggering access, pay, routing, or deactivation consequences without a sufficiently legible review layer	Separated review authority: human review node between quality signal and irreversible consequence
Memory	Per-project only; contributor resets to zero at each project boundary	Portable contributor continuity record; state persists and governs routing and consequence

New system objects

Reliability stream gate: Qualifies contributors for Stream A by testing actual execution capability, not credential density

Expert judgment gate: Qualifies contributors for Stream B by testing specialist reasoning and reframing capability

Portable continuity record: Preserves contributor state across projects; governs routing, consequence, and review

Consequence review node: Human review layer between algorithmic quality signal and irreversible action

Expert signal channel: Governed route for Stream B contributors to surface rubric and task-definition problems to a layer with authority to act on them

Queue-state visibility layer: Surfaces contributor-relevant queue-state reasons where they govern access, inactivity, or routing status

Formal interpretation layer: Bounded, governed interpretation function replacing the informal QM interpretation burden; carries basic system meaning so people do not have to

SECTION 5

MINIMUM VALID REPLACEMENT

The redesign is only real if all five of the following exist simultaneously. A deployment missing any one of these is not a replacement. It is a partial correction inside the old frame.

- → **1. Separate gates** for reliability execution and expert judgment — different qualification logic testing different capabilities, not one gate with two labels
- → **2. Portable contributor continuity record** that persists across projects, is visible to contributors, and actively governs routing and consequence decisions
- → **3. Visible queue-state reasons** surfaced in the platform interface without requiring Queue Manager explanation or external tooling
- → **4. Human review before irreversible consequence** between algorithmic quality signal and deactivation or pay reduction
- → **5. Expert signal channel that changes upstream behaviour** — not just receives submissions; actually alters task, rubric, or boundary definitions

MINIMUM REPLACEMENT RULE

If any one of these five components is missing, the system is still correcting inside the old frame. The minimum replacement is not a starting point for negotiation. It is the threshold below which redesign has not begun.

The full Redesign Artifact expands these five replacement components into a seven-part minimum valid pilot by adding pre/post burden measurement and pilot-specific validation rules. The five here define the architecture. The seven define the pilot.

Board decision required

BOARD QUESTION	REQUIRED ANSWER
What work does the platform actually buy?	Reliability execution, expert judgment, or both — as separate streams with separate gates, metrics, and compensation. Not one frame claiming both.
What must stop before redesign?	Credential-density intake, hidden queue-state logic, one-frame routing, unpaid onboarding as a coherence buffer, volume scaling

BOARD QUESTION	REQUIRED ANSWER
	before gate redesign.
What must be built?	Separate stream gates, portable continuity record, queue-state visibility, consequence review node, expert signal channel that alters upstream behaviour.
What proves success?	Same pressure produces different response: QM burden falls, continuity governs routing, expert signals change task definitions, high-signal churn declines.
What invalidates success?	Stream B is symbolic, compensation converges to throughput rates, QMs remain load-bearing, automated scoring still triggers irreversible consequence without review.

SECTION 6

WHAT STOPS IMMEDIATELY

Before the redesign can hold, the following must stop. These are preconditions, not optimisation targets. Redesign attempted while any of these continue will replicate the mismatch in its first phase.

STOP	WHY
Using credential density as a proxy for production quality	Credentials predict expertise. The system is not primarily designed to consume expertise. The gate selects the wrong population.
Treating longer guidelines as the primary response to quality variance	More pages intensify the signal-switch. They convert expert labour into compliance labour earlier and more completely.
Treating the Empty Queue as a communication problem	Queue-state reason codes are an information-architecture condition, not a communication problem. When state governs access or inactivity and is not surfaced, contributors and Queue Managers carry the interpretation burden the platform should hold structurally.
Scaling contributor volume before primitive correction	Every additional credentialed contributor through the old gate compounds recruitment waste before the redesign can hold.
Routing reliability work and expert judgment through the same gate	A single gate cannot test for two incompatible work classes. It optimises for one and systematically misclassifies the other.
Using unpaid onboarding as a system coherence buffer	Unpaid onboarding hides the cost of guideline complexity and project volatility. The redesigned system must not inherit this burden-transfer model.

SECTION 7

DEPENDENCY ORDER

The redesign must happen in the correct sequence. Right fixes in the wrong order recreate the failure at a new stage. The current failure pattern persists precisely because interface-level interventions — Move 3 and Move 4 actions — are attempted without Move 1 in place.

MOVE	WHAT HAPPENS	CONSEQUENCE IF SKIPPED
1	Name the primitive correctly at leadership level — formally declare the two incompatible work classes	All downstream changes are interpreted through the old frame and replicate the mismatch

MOVE	WHAT HAPPENS	CONSEQUENCE IF SKIPPED
	and retire the single broken frame	
2	Redesign qualification gates into stream-specific logic testing actual stream capability	Stream split is cosmetic; both streams admit the wrong population
3	Build portable contributor continuity record	Contributors still reset to zero; routing cannot improve despite stream separation
4	Install queue-state visibility, then consequence review — first surface what the queue state is; then separate scoring from irreversible consequence	Hidden burden remains load-bearing; legal exposure pathway remains open
5	Install expert signal channel with authority to change upstream task behaviour	Stream B is a prestige label for reliability work; recapture is immediate
6	Retire the single-frame architecture	Old frame quietly reasserts itself under volume pressure

The dependency rule: each move must be complete before the next move scales. Preparatory work may begin early, but no component may become the official correction path before its dependency is in place. Attempting Move 3 or Move 4 without Move 1 in place is exactly why every previous correction cycle has ended in the same place.

SECTION 8

FAKE PROGRESS TEST

The redesign can be implemented in appearance while the old failure continues in structure. The following conditions confirm that the primitive has returned, not that the redesign has succeeded:

Two streams exist in language but share one qualification gate
Not a redesign. Two labels on the same broken frame.

The continuity record exists but does not affect routing or consequence decisions
The record is a document, not a system object. It must govern decisions.

Queue-state visibility exists but Queue Managers still have to explain it
The visibility layer is not functioning. QMs remain load-bearing.

Expert contributors can submit signals but nothing changes upstream
The signal channel is a filing system. Stream B is a label.

Automated scoring still triggers irreversible consequence without human review
The sensing-authority fusion is intact. Legal exposure is unchanged.

Stream B contributors are compensated at throughput rates
Compensation recapture. The stream split is cosmetic regardless of gate design. Directly violates the stream contract (Section 3).

Do not approve a redesign plan that:

- → Creates two stream labels but keeps one qualification gate
- → Exposes queue-state status but leaves Queue Managers explaining what it means
- → Creates an expert stream but gives it no authority to change task definitions
- → Creates a continuity record that does not affect routing or consequence decisions
- → Pays expert judgment at throughput labour rates
- → Leaves automated score-to-consequence pathways intact without a human review node

SECTION 9

CONSEQUENCE OF DEFERRAL

The current correction loops all operate inside the malformed frame. Each cost class below compounds without primitive correction:

> ***COST CLASSES — FOUR STRUCTURAL CONSEQUENCES OF DEFERRAL***

- → **Asset depreciation:** every credentialed specialist onboarded into the current frame is a high-cost asset depreciated toward zero within one project cycle.
- → **Shadow payroll:** Queue Managers, senior contributors, and community tooling carry coherence functions the platform does not formally own.
- → **Legal tail-risk:** the sensing-authority fusion is structurally consistent with active legal scrutiny; crystallisation timing is uncertain.
- → **Upstream signal loss:** compliance-pattern data may be delivered where specialist judgment is expected or implied by the recruitment surface. This cost does not appear in internal QA metrics because the system has no sensor for it.

Recruitment waste, support overhead, quality-control inflation, contributor churn, and legal exposure will continue to compound. The compounding dynamic has no internal exit without primitive correction:

STAGE	MECHANISM	STRUCTURAL ORIGIN
1	Recruitment waste increases cost pressure	Wrong gate admits wrong population
2	Cost pressure → more algorithmic quality management	No primitive correction; execution-layer response
3	More QA → more scoring opacity	Sensing-authority fusion deepens
4	More opacity → more contributor frustration	No queue-state visibility; no appeal mechanism
5	More frustration → more complaints and legal action	Legal and reputational exposure increases
6	More compliance → stricter rubrics → wider expert-compliance gap	Guidelines deepen the signal-switch
7	Churn rises in high-signal cohorts	High-signal contributors are more likely to churn or disengage when specialist value is not retained
8	More recruitment waste → return to Stage 1	No internal exit without primitive correction

The deferral rule: *any intervention that leaves the primitive intact will reproduce the same failure pattern in a new form.*

SECTION 10

PROOF CONDITIONS

The redesign is real when it passes the same-pressure test:

> **PROOF STANDARD**
> ***Same pressure → different response.***

In one sentence for the board: *when the platform stops punishing expertise and starts routing it to work that can use it, the redesign is real.*

Before redesign begins: baseline measurements must be captured — contributor churn by tier, QM interpretation ticket volume, Empty Queue duration, deactivations without human review. Without a baseline, different-response cannot be demonstrated.

The redesign is confirmed when all of the following hold:

- → Contributors do not reset to zero across projects — continuity is live and governs routing
- → Queue state is visible in the platform without external tools or QM explanation
- → QM interpretive support load has measurably fallen
- → Automated scores cannot trigger irreversible consequence without Review Node approval
- → **The expert signal channel has changed upstream behaviour:** at least one rubric, task boundary, or review rule has been altered by a Stream B signal during the pilot. A channel that receives submissions but alters nothing is a filing system, not a signal channel.
- → High-signal contributor churn has measurably declined

SECTION 11

BOARD VERDICT

BOARD VERDICT

Replace the single broken contributor frame with a dual-stream architecture that separates reliability execution from expert judgment, restores contributor continuity across projects, and removes the hidden coherence burden currently carried by Queue Managers and contributors. The stream split is fake without stream-specific compensation: expert judgment cannot be paid as throughput labour regardless of gate design. The redesign is not an optimisation of the current system. It is a structural replacement of the primitive that generates the failure.

The board choice: is not between redesign and optimisation. Optimisation has already failed because it acts inside the malformed frame. Every interface fix, dashboard update, and support escalation has operated on a symptom while the cause regenerated. The choice is whether to correct the primitive or continue compounding the costs it produces.

The full Redesign Artifact is restricted and contains the complete transition sequence, move-by-move logic, rules and invariants, drift immune system with fail-closed primitive regression gates, admissibility conditions and stall-risk mitigations, pilot validity rules, object-level builder handoff, and automated testability requirements.

Next steps for the board

- → **1. Accept the governing diagnosis:** sign the Diagnostic Acceptance Protocol (Diagnostic Artifact, Section 15). 14-day deadline. This acknowledges the primitive mismatch, stops the Vanish List items, and is the formal trigger for the Redesign Artifact engagement.
- → **2. Commission the full Redesign Artifact:** the restricted builder specification contains the complete transition sequence, invariants, drift gates, admissibility conditions, and builder handoff. It is released after the Diagnostic Acceptance Protocol is signed.
- → **3. Contact:** hello@jamieforrester.com · jamieforrester.com to discuss the engagement.

WHAT THIS DEMONSTRATES

This executive summary demonstrates what a structural replacement specification produces and how this practice approaches primitive-level systems failure. The full Redesign Artifact sits underneath it with complete builder specification and governed transition logic. The diagnostic and redesign together form a governed artifact pair — the diagnosis proves the failure; the redesign specifies the replacement. This public summary is not sufficient for implementation. It demonstrates the replacement logic while withholding the restricted transition sequence, invariants, drift gates, admissibility work, and builder-level architecture.

Delivery model: this redesign is specified independently of delivery model. The dual-stream architecture and contributor continuity record must remain under the platform's operational control and may not be transferred to a third-party vendor in a way that creates a new form of capture or custody loss.

To discuss access to the restricted Redesign Artifact or how this diagnostic and redesign methodology applies to your organisation, contact:

hello@jamieforrester.com · jamieforrester.com

Prepared by Jamie Forrester, Independent Systems Architect, Edinburgh, UK. Produced independently — not commissioned by Outlier or Scale AI. Derived from the Diagnostic Artifact (public) and the full Redesign Artifact (restricted). This summary introduces no independent evidentiary claims. It summarises the replacement logic derived from those two artifacts.