PUBLIC TRANSLATION ARTIFACT · 01 IN READING ORDER

OUTLIER

/ AI · PLATFORM RELIABILITY

PLAIN-LANGUAGE COMPANION

For leadership, product teams, clients, and builders.

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

ARTIFACT FAMILY

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

PREPARED BY Jamie Forrester · Independent Systems Architect · Edinburgh, UK
CONTACT hello@jamieforrester.com · jamieforrester.com
DATE April 2026
CLASSIFICATION Public · 01 in reading order · start here

QUICK REFERENCE

THE PROBLEM	THE FIX
Outlier recruits through an expertise market but runs a reliability pipeline. One broken frame coordinates two incompatible work classes.	Replace the single frame with a dual-stream architecture. Stream A for reliability. Stream B for expert judgment. Both connected by a portable contributor continuity record.

THE 6 MOVES	PLAIN-LANGUAGE PROOF SIGNALS
1. Declare the corrected primitive	Contributors do not reset to zero across projects
2. Redesign the qualification gates	Queue state is visible without external tools
3. Build contributor continuity	Queue Managers are no longer daily coherence carriers
4. Expose queue state + review consequence	Automated scores cannot trigger irreversible consequence
5. Install the expert signal channel	Expert signals can lead to task-definition and rubric changes
6. Retire the single-frame system	High-signal contributors stay instead of churning

Plain-language terms used in this document

TERM	PLAIN MEANING
Stream A	The reliability execution track. Contributors perform bounded, rubric-defined tasks at volume. Measured on consistency and throughput.
Stream B	The expert judgment track. Credentialed specialists perform work requiring domain reasoning, rubric challenge, and signal generation. Measured on quality and upstream impact.
Queue Manager (QM)	A role currently carrying coherence work the platform architecture should do itself — explaining queue state, routing contributors, absorbing system ambiguity. In the redesign, that burden moves into structure.
Contributor continuity record	A portable record of a contributor's work history, stream, performance, and qualifications that follows them across projects. Currently absent; contributors reset to zero with each project.
Ghost structure	The hidden human work — QMs explaining queue logic, senior contributors bridging ambiguity, contributors maintaining their own tracking spreadsheets — that the architecture should be doing but is not.
Collapse node	The point where too many system failures converge at once. Here: the single-frame architecture that compresses two incompatible work types into one coordination model.
Recapture	The old failure returning inside the new system. When Stream B begins to be measured on Stream A criteria, recapture has occurred.
Vanish List	The non-negotiable list of things that must stop before the redesign can hold. Items on the Vanish List cannot be negotiated, phased, or deferred.

SECTION 0

WHAT THIS DOCUMENT IS

This plain-language document explains what went wrong at Outlier, why previous attempts at correction have not worked, and what a real replacement requires — without requiring you to read the full Diagnostic or Redesign Artifacts first. Use it to get everyone on the same page before the full artifacts are circulated.

It serves as a companion to the full artifacts:

- → **Diagnostic Artifact —** proves the structural failure with full evidence and evidence tiering
- → **Redesign Artifact —** specifies the dependency-ordered replacement architecture for builders
- → **Redesign Executive Summary —** public board summary of the replacement logic
- → **Translation Artifact —** this document; creates shared understanding across leadership, teams, and stakeholders

START HERE — READING ORDER

Recommended reading order: Start here (Translation Artifact) → Diagnostic Artifact (full evidence) → Redesign Executive Summary (replacement logic) → Full Redesign Artifact (restricted — not publicly published; contains complete builder specification).

Its job is translation: to make the diagnosis and redesign understandable before the full artifacts are circulated. It preserves the judgment boundary of the Diagnostic Artifact: this is a structural reading from available evidence, not a claim about undisclosed internal systems or intent.

Who this is for

This document is for everyone who needs to understand the diagnosis and redesign without reading the full technical specification first.

AUDIENCE	WHAT TO TAKE FROM THIS DOCUMENT	THEN READ
Leadership and board	Understand the primitive failure, why it is expensive, and why the redesign sequence matters. Accept or reject the governing diagnosis.	Diagnostic Artifact (full evidence)
Product and operations	Understand what stream separation means, what must stop, and what the minimum redesign requires.	Redesign Artifact (Sections 5–12)
Engineering	Understand the required system objects: continuity, visibility, consequence review, signal channel.	Redesign Artifact (Sections 4, 6, 16)
Queue Managers and operations teams	Understand how burden must move out of people into structure, and what that means for your current role.	Redesign Artifact (Section 7)
External stakeholders and reviewers	Understand the structural diagnosis without needing the technical specification.	Diagnostic Artifact (public)

What this document does not say

Readers may be tempted to draw conclusions the document does not support. To be precise:

- → This document does not claim Outlier or Scale AI acted in bad faith or with intent to deceive
- → This document does not say algorithmic scoring must be removed — it says scoring must not directly trigger irreversible consequence without human review
- → This document does not say Queue Managers should be removed — it says their role must change from load-bearing coherence to supplementary escalation
- → This document does not say all work should become expert judgment work
- → This document does not say Outlier must guarantee continuous work availability
- → This document does not claim access to internal systems, dashboards, or client contracts

What it does say: the current contributor frame collapses two incompatible kinds of work into one system, and that collapse creates a recurring structural failure that surface corrections cannot resolve.

SECTION 1

THE SYSTEM IN ONE SENTENCE

> GOVERNING DIAGNOSIS
>
> ***Outlier uses a single contributor frame to coordinate two incompatible kinds of work — reliability execution and expert judgment — so it recruits for expertise but measures for compliance.***

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

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

SECTION 2

WHAT THE SYSTEM SAYS VS WHAT IT DOES

The platform signals expertise at entry: specialist judgment, high-skill contribution, and domain-specific capability. In production, it measures something different: rubric adherence, time-bound outputs, formatting compliance, and variance suppression.

WHAT THE PLATFORM SIGNALS AT ENTRY	WHAT THE PLATFORM MEASURES IN PRODUCTION
Specialist judgment and domain expertise	Rubric adherence and formatting compliance
High-skill contribution	Time-bound bounded task completion
Expert reasoning matters	Variance suppression — correct variance is penalised, not valued
Your credentials qualify you for this work	Your compliance score determines your access

This is not a contributor problem. It is a structural mismatch between what the system signals and what it consumes. Where a contributor's specialist knowledge leads them to diverge from the rubric, the visible scoring architecture appears to reward rubric conformity over correct domain variance. The result is not that expertise is absent. The result is that expertise becomes operationally unsafe when it conflicts with the rubric. The mismatch is embedded in the platform's design, not in the quality of the people it recruits.

SECTION 3

WHERE THE BURDEN LANDS

The platform currently functions — tasks are completed, scores are issued, projects proceed — because people outside the formal architecture perform the work the architecture does not carry.

Support helps a working system run. Coherence labour makes an incoherent system usable. In Outlier's current structure, too much coherence labour is carried by people outside the formal design:

Queue Managers

Burden: absorb the interpretation of shifting guidelines and queue opacity

For example: a Queue Manager may spend an hour in a webinar or Slack thread explaining why a qualified contributor's queue is empty — because the contributor-facing interface does not surface the reason. That is not support work. That is structural incoherence carried by a person.

Contributors

Burden: carry unpaid onboarding cycles, repeated requalification, queue-state inference from external forums, and guideline interpretation not compensated as task time

Contributor reports describe hundreds of hours spent on processes the platform does not formally recognise as work

External tools and community forums

Burden: surface information available in the platform environment but not surfaced in the contributor-facing interface

The EmptyQueue-Extension — a third-party GitHub tool — was built specifically to surface queue-state reason codes available in the platform environment but not surfaced to contributors. When open-source tooling exists to correct a platform's own information architecture, the structural failure has exceeded its own carrier layer.

> STRUCTURAL READING
>
> ***These are not support roles. They are a load-bearing coherence layer. Remove them and the platform becomes operationally incoherent. The redesign succeeds only when this burden moves from people into structure.***

SECTION 4

COST OF CONTINUATION

The current structure does not produce a stable, manageable cost. It produces a compounding one.

ASSET DEPRECIATION

Every credentialed specialist recruited into the wrong stream risks becoming a high-cost asset whose specialist value is not retained beyond one project cycle. The platform spends intake cost to acquire specialist capability, then loses that capability by placing it inside a system that treats specialist variance as failure.

The cost classes

COST CLASS	PLAIN-LANGUAGE DESCRIPTION	DIRECTION UNDER DEFERRAL
Recruitment waste	Every expert who churns after entering the wrong stream represents full intake cost with zero retained value	Compounds — more recruitment to replace churn
Support overhead	Every hour a Queue Manager interprets queue state is cost created by the system's own opacity — not by the work itself	Grows — QM infrastructure is now permanent, not temporary
QA inflation	Every compliance layer added to fix quality variance increases cost without reducing the variance, because the variance originates in the primitive not the execution	Non-linear — each layer adds cost without fixing the source
Shadow payroll	Contributor inference, onboarding, and guideline interpretation work that does not appear in operating budgets but generates legal tail-risk	Invisible — unrecognised until it crystallises as legal action
Legal and regulatory exposure	Multiple lawsuits and PAGA claims, a reported DOL investigation, and investigative journalism consistent with the gap between what the platform signals and what it does	Tail-risk severe; timing uncertain
Reputational damage	Rising contributor acquisition friction; accelerating press coverage; 1/10 Oxford Internet Institute labour standards score	Directional and observable

The compounding dynamic: recruitment waste increases cost pressure → more algorithmic quality management → more scoring opacity → more contributor frustration → more complaints and legal action → stricter compliance → wider expert-compliance gap → more churn → more recruitment waste. There is no internal exit from this loop without correcting the primitive.

Why this matters upstream: the risk is not only contributor frustration. If a platform recruits expert contributors but operationally rewards rubric conformity, the signal delivered upstream may differ from the capability implied by the recruitment surface. Compliance-pattern output may scale where specialist judgment was expected or implied. This is a structural risk, not a claim about any specific client contract, model outcome, or internal quality standard.

SECTION 5

WHY PREVIOUS FIXES HAVE NOT WORKED

Outlier has tried stricter credential requirements, longer and more detailed guidelines, dashboard redesigns, Marketplace navigation improvements, Queue Manager support layers, and webinars. None have resolved the core failure.

> **THE KEY DISTINCTION**
>
> ***Interface correction improves performance inside the current frame. Redesign changes the frame itself. Every previous fix has been an interface correction. None have reached the frame.***

A specific example: longer guidelines were introduced in response to quality variance. The logic was that if contributors understood the rubric better, compliance would improve. What happened instead: longer guidelines intensified the signal-switch. Expert contributors were required to spend more time learning to suppress their specialist judgment. The guidelines made the mismatch deeper, not shallower.

The pattern holds for every fix attempted:

FIX ATTEMPTED	WHY IT COULD NOT WORK
Stricter credential requirements	Credentials predict expertise. The system is not designed to consume expertise. A stricter gate selects more qualified people for the wrong work class.
Longer guidelines	More pages intensify the rubric-compliance requirement. They deepen the expert-to-compliance conversion, not correct it.
Dashboard and interface improvements	Interface changes cannot correct a structural primitive. Making the Empty Queue more visible does not correct the mismatch that generates it.
More Queue Manager support	QM support absorbs symptoms. It cannot correct the source of the ambiguity it is interpreting.
Scaling volume	More contributors through a broken gate produce more failure of the same kind, not different failure.

As long as the primitive remains mismatched, every new correction layer simply relocates the same failure to a new surface.

SECTION 6

THE REDESIGN IN PLAIN LANGUAGE

The redesign replaces the broken single-frame structure with a dual-stream architecture. This is not an incremental improvement. It is a structural replacement of the primitive.

	STREAM A — RELIABILITY EXECUTION	STREAM B — EXPERT JUDGMENT
What this stream is for	Repeatable, constrained work where rubric conformance is the product	Specialist reasoning where correct variance is the product
Who should enter it	People qualified for reliable, bounded execution	Domain specialists whose divergence from a rubric may be the most valuable output
How success is measured	Reliability, rubric conformance, throughput within realistic time limits	Signal quality, rubric challenge quality, upstream impact of corrections submitted

	STREAM A — RELIABILITY EXECUTION	STREAM B — EXPERT JUDGMENT
What variance means	Suppressed — variance is the failure signal	Valued — correct variance is what the stream is purchasing
Compensation logic	Per-task or per-time aligned to execution work; paid onboarding where required	Per judgment contribution; reflects specialist input value, not volume compliance

Stream contracts

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

If Stream B exists but contributors in it are paid and scored like Stream A contributors, the stream separation is cosmetic. The mismatch continues under a different name.

The continuity record

Both streams share a portable contributor continuity record — a permanent record that follows the contributor across projects. It preserves what they have qualified for, where they have performed well, what consequences have occurred, what disputes have been resolved, and what routing decisions should consider before assigning future work.

The continuity record must be visible enough to be contested. It cannot become another hidden file. A contributor should be able to see what the system knows about them and request formal review of any entry. If the continuity record is invisible to contributors, it recreates the same structural opacity as the current queue state.

The expert signal channel

The most important element of Stream B is the expert signal channel: a governed route through which expert contributors can surface problems with rubrics, task definitions, or boundary conditions to a layer with actual authority to act on them.

Concrete example: an expert contributor identifies that the rubric for a mathematics task contains an error that causes systematically worse outputs. They submit a correction through the signal channel. The channel routes it to a review node with authority to update the rubric. The rubric is corrected. The fix propagates across the task type. That is what the expert stream is for. A signal channel that receives submissions but changes nothing is just a label.

SECTION 6.1

WHAT THE MINIMUM REDESIGN REQUIRES

The redesign is only real if five things exist together. 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.** Not one gate with two labels. Different qualification logic testing different capabilities.

2. **A portable contributor continuity record** that follows contributors across projects, is visible to them, and actively governs routing and consequence decisions.
3. **Visible queue-state reasons.** Contributors can see — in the platform, without external tools — why they have or do not have available work.
4. **Human review before irreversible consequence.** An algorithmic score can flag risk, but cannot alone remove access or reduce pay. A human review node must sit between signal and action.
5. **An expert signal channel that can change task, rubric, or boundary behaviour.** Not just receive submissions. Actually change outcomes.

MINIMUM REPLACEMENT RULE

If any of these five are missing, the change is not a redesign. It is a partial correction inside the old frame — and the old failure will return.

SECTION 7

THE MOVE ORDER

The redesign must happen in sequence. Right fixes in the wrong order recreate the failure.

MOVE	WHAT HAPPENS	WHY IT MUST COME HERE
1	Declare the primitive — leadership clearly names the two different work classes and retires the single broken frame	All downstream design depends on this. Without it, every subsequent change is interpreted through the old frame.
2	Redesign the gates — reliability and expert judgment receive separate qualification logic testing actual stream capability	Without the correct gate, the stream split is cosmetic. Both streams admit the wrong people.
3	Build contributor continuity — history, fit, and routing relevance persist across projects	Without continuity, contributors still reset to zero despite stream separation. Routing cannot improve.
4	Expose queue state and review consequence — contributors see why they have or do not have work; automated scores cannot directly trigger irreversible actions	Without visibility, the hidden burden remains. Without review separation, the legal exposure pathway remains open.
5	Install the expert signal channel — expert contributors can surface malformed rubrics or task problems to a layer with authority to act	Without this, Stream B is a prestige label for reliability work. The separation is symbolic.
6	Retire the single-frame system — the old unified contributor architecture is removed from active use	Without retirement, the old frame quietly reasserts itself under pressure.

The sequence is part of the redesign. Skipping steps does not accelerate correction; it moves the old failure into the new structure at a different stage.

SECTION 8

WHAT MUST STOP

Before the redesign can hold, six behaviours must stop. These are not optimisation targets. They are preconditions. In this document, irreversible consequence means any action that materially affects contributor access, pay, tier status, project eligibility, or account standing and cannot be automatically reversed without formal review. The redesign attempted while any of these continue will replicate the mismatch in its first phase.

STOP	WHY
Using credentials 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 main response to quality variance	Longer guidelines intensify the signal-switch. They convert expert labour into compliance labour earlier and more completely. More pages deepens the mismatch.
Not surfacing queue-state	The Diagnostic Artifact identifies queue-state reason codes not surfaced

STOP	WHY
logic to contributors	to contributors. Where those states govern access, they should be translated into contributor-facing explanations. Hidden state must be carried by people.
Scaling contributor volume through the broken single frame	Every additional credentialed contributor through the old gate deepens the mismatch 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 will optimise for one and misclassify the other.
Using unpaid onboarding as a system coherence buffer	Unpaid onboarding hides the cost of guideline complexity. The redesigned system must not inherit this burden-transfer model.

SECTION 9

WHAT WOULD COUNT AS FAKE PROGRESS

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

The two streams exist in language but share one qualification gate
Not a redesign. It is two labels on the same broken frame.

The continuity record exists but does not affect routing decisions or contributor access
The record is a document, not a system object. The redesign requires it to govern decisions.

Queue-state visibility exists but Queue Managers still have to explain what it means every day
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.

Score-mediated signals remain capable of triggering irreversible consequence without a sufficiently legible human review layer
The sensing-authority fusion is structurally intact. The structural exposure pathway remains.

Stream B contributors are paid and scored like throughput workers
Compensation recapture. The stream split is cosmetic regardless of what the gates say.

The redesign is declared complete before baseline measurements are taken
Without a baseline, same-pressure / different-response cannot be demonstrated.

These are not minor implementation gaps. They are signs that the old primitive has returned under a new name. The test is always: has the burden actually moved from people into structure?

SECTION 10

WHAT THE SYSTEM MUST NOW CARRY ITSELF

This is the redesign in its simplest form: everything currently carried by people as workaround labour must become an explicit system function. The redesign succeeds only when hidden burdens move from people into structure.

CURRENT BURDEN CARRIED BY PEOPLE	MUST BECOME	PLAIN-LANGUAGE MEANING
Contributor history resets at each project	Persistent continuity record	The system remembers what contributors have already proven, so they do not start from zero each time
Queue uncertainty — contributors must infer why they have no work	Visible queue-state reasons in the platform	Contributors can see why they do or do not have work, without asking anyone or using external tools
Guideline interpretation by Queue Managers	Formal interpretation layer in the platform architecture	Basic system meaning is carried by the platform, not by informal human bridges
Score-mediated signals appear capable of triggering access, pay, or deactivation consequences without a sufficiently legible review layer	Reviewed decision path — human review before irreversible consequence	A score can flag a risk, but cannot alone trigger access, pay, or deactivation consequences without a human review node
Expert insight submitted but ignored	Expert signal channel that can change task and rubric behaviour	Specialist contributors can challenge flawed rubrics, and those challenges can actually change the task
Stream promise mismatch — recruited as expert, measured as throughput	Stream contracts — measured against the work you were recruited to do	Reliability work is not dressed as expert judgment; expert judgment is not paid as throughput
Unpaid onboarding and repeated requalification	Paid readiness where required by the platform	If the platform requires preparation for production, that preparation is recognised as work

If contributors or Queue Managers still need to infer, translate, or compensate manually after the redesign, the burden has not moved. The redesign has not succeeded.

SECTION 11

WHAT SUCCESS LOOKS LIKE

When the redesign is working, the experience of the platform changes concretely.

Before and after

Before: a credentialed specialist joins a new project. Their work history from the previous project does not carry forward. They requalify from scratch. Their queue is empty and no reason is shown; they check community forums and ask their Queue Manager. They submit a correction to a rubric error they have identified. Nothing changes. They churn after one project cycle.

After: the same specialist enters through Stream B. Their continuity record from the previous project is live — their performance, qualifications, and routing strengths carry forward. Their queue state is visible and explained. When they identify the rubric error, they submit it through the expert signal channel. Two weeks later, the rubric has been updated and the fix propagates across the task type. They stay.

For a contributor

A specialist who qualifies for Stream B enters a project and immediately sees their queue state — what is available, what is not, and why. When they complete work, their performance is preserved and carries forward to the next project. They do not start from zero. When they identify a structural problem with a rubric, they submit it through the signal channel. Two weeks later, the rubric has been updated. Their contribution changed the system.

For a Queue Manager

The volume of contributors asking "why is my queue empty?" has fallen materially because the platform now shows the answer. The webinars about queue logic are no longer needed for basic orientation. QMs spend their time on genuine escalations — unusual situations that warrant human judgment — rather than translating the platform's own ambiguity back to contributors every day.

For leadership and the platform

Credentialed specialists who enter Stream B stay longer because the work matches the contract they were offered. High-signal contributor churn among expert cohorts has fallen. The structural source of legal and regulatory pressure is reduced because the promise and the delivery now align. The platform's recruitment language and its operating mechanics say the same thing.

> SUCCESS CONDITION
>
> ***The redesign is not complete when the architecture has changed. It is complete when the experience has changed. Burden has moved from people into structure. That is the test.***

SECTION 12

WHAT PROVES THE REDESIGN IS REAL

The proof condition is simple:

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

Before redesign begins: baseline measurements must be captured. Without them, different-response cannot be demonstrated. Measure: contributor churn by tier; Queue Manager interpretation ticket volume; average Empty Queue duration; deactivations without human review; repeat requalification events; external tool usage rates.

The redesign is real when all of the following hold:

- → Contributors do not reset to zero across projects — continuity is live and governs routing
- → **Queue state is visible without external tools or QM explanation**
- → Queue Managers are no longer daily coherence carriers — their load has measurably fallen
- → Automated scores cannot directly trigger irreversible consequence without human review
- → Expert contributors can influence task definitions — the signal channel produces upstream changes
- → High-signal contributors stay instead of churning early

The redesign is not proven if: the old burden still needs a human workaround layer — if continuity does not govern decisions, if queue state is visible but QMs still have to interpret it, or if expert signals produce no upstream change. These are signs that the old primitive has returned inside the new architecture.

SECTION 13

RECAPTURE WARNING

The redesign's greatest enemy is not initial resistance. It is quiet recollapse after apparent success. As volume and pressure increase, the system will face forces that push toward the old structure. These forces must be anticipated and structurally blocked.

RECAPTURE RISK	HOW IT HAPPENS	DETECTION SIGNAL
Stream collapse	Pressure to merge both streams into one system for speed or simplicity — reinstalls the original failure	One gate or routing model begins serving both work classes again
New hidden carriers	Queue Managers or senior contributors again become the informal interpretation bridge	Support and forum dependence rise instead of fall despite redesign rollout
Governance recentralisation	Stream-specific authority collapses back into throughput management under pressure	Review layer is bypassed; stream-specific decisions disappear
Audit memory replacing operational continuity	The continuity record grows for legal defensibility but stops affecting routing or contributor protection	Continuity record exists but does not govern actual system decisions
Expert stream becoming a label	Expert stream exists in language only, without authority to affect task definitions	Expert contributors can submit signals but nothing upstream changes
Compensation recapture	Both streams are paid as throughput labour despite stream separation	Expert stream compensation rates collapse to Stream A rates; expert contributors exit

RECAPTURE PRINCIPLE

The dual-stream structure is not a preference. It is the corrected primitive. If the gates or metrics merge, the mismatch returns. If the expert signal channel becomes just a label, the expert stream becomes cosmetic. The redesign only survives pressure if the separation survives pressure.

SECTION 14

HOW TO USE THIS DOCUMENT

Use this Translation Artifact before the full artifacts, not instead of them.

AUDIENCE	HOW TO USE THIS DOCUMENT	NEXT STEP
Leadership and board	Use it to understand the primitive failure, why it is expensive, and why the redesign sequence cannot be skipped or cherry-picked	Accept the governing diagnosis formally. Sign the Diagnostic Acceptance Protocol. Declare Move 1.
Product	Use it to understand stream separation, stream contracts, and what fake progress looks like	Read Redesign Artifact Sections 5–6 for specification; Section 8 for the Vanish List
Engineering	Use it to understand the required system objects and proof conditions	Read Redesign Artifact Sections 4, 6, 16, and 19 for build specification and testability
Operations / QM layer	Use it to understand how the QM role must change: from load-bearing coherence to supplementary escalation	Read Redesign Artifact Section 7 for the burden carrier map
Legal / compliance reviewers	Use it to understand the structural origin of the consequence exposure	Read Diagnostic Artifact Sections 7 and 11 for the sensing-authority fusion and cost of continuation
External stakeholders	Use it as the accessible entry point to the Outlier structural diagnosis	The Diagnostic Artifact (public) is the full evidence base

SECTION 15

WHAT TO DO NEXT

For leadership and boards reading this document:

- → **1. Accept the governing diagnosis.** Outlier is an expertise-coded reliability pipeline. The failure is structural, not executional. This acknowledgment is the prerequisite for every subsequent move.
- → **2. Stop the Vanish List items.** Credential-density proxy, longer guidelines as a fix, invisible queue-state logic, volume growth before primitive correction, one-frame routing, unpaid onboarding as coherence buffer. These stop before Move 1.
- → **3. Sign the Diagnostic Acceptance Protocol.** Binding, 14-day deadline. This formally acknowledges the primitive mismatch and is required before the full Redesign Artifact is released. Without a signed acceptance, the redesign engagement does not proceed.
- → **4. Commission the Redesign Artifact.** The restricted builder specification contains the move-by-move transition sequence, rules and invariants, drift immune system, admissibility conditions, and builder handoff.
- → **5. Start Move 1.** Primitive declaration at leadership level. This is not a dashboard, gate, or workflow change. It is a formal statement that reliability execution and expert judgment are different work classes and must be coordinated differently.

CLOSING STATEMENT

This document has named the failure, explained why previous fixes could not reach

it, and described what real replacement requires. The diagnosis is not commentary. The redesign is not a recommendation. They are structural findings and a structural replacement specification. In a live engagement, redesign begins only after the governing diagnosis is formally accepted. The first move is a leadership-level primitive declaration: reliability execution and expert judgment are different work classes and must be coordinated differently.

WHAT THIS DEMONSTRATES

This Translation Artifact demonstrates that a structural diagnosis of this kind is explainable across audiences without weakening the governing finding. The diagnostic logic, the replacement architecture, and the recapture resistance conditions are intact. The public artifacts together form a proof of method: the same structural framework applies across a private AI platform and a public healthcare institution. The practical question is whether the primitive is corrected before the same failure pattern reproduces through another correction cycle.

For the full evidence base, see the Diagnostic Artifact (Artifact 01, public). For the builder specification and dependency-ordered transition sequence, see the Redesign Artifact (Artifact 02, restricted). To discuss the engagement, contact: hello@jamieforrester.com · jamieforrester.com

Prepared by Jamie Forrester, Independent Systems Architect, Edinburgh, UK.