

AN INTELIIUM WHITEPAPER

# Defensible Flow

*AI is rebuilding the software organization, not just the toolchain.*

---

**Andrew Rufener**, Intelium Consulting  
**Angus Gow**, Flowfoundry

## CONTENTS

# Defensible Flow

---

Executive Summary .....	3
1. The Conversation So Far, and the Missing Map.....	4
2. The Constraint, Sharpened: Defensible Flow.....	5
3. The Technical Shape of the Target.....	6
4. The Organizational Change — the Hard Part .....	8
4.1 The gates are people .....	8
4.2 The seniority inversion .....	9
4.3 The coordination fabric becomes an organizational design problem.....	10
4.4 Measures and incentives move, or nothing else does .....	10
4.5 The cultural preconditions are testable .....	11
5. Getting There: a Staged Hypothesis.....	11
Stage 1 — Build the fabric before scaling the agents .....	11
Stage 2 — Run the new model inside the old governance.....	12
Stage 3 — Retire gates by decision class, by evidence.....	12
6. What Leaders Should Do Now.....	13
Where are your decisions, actually? .....	13
Could you produce the file? .....	13
Do your gate-operating functions have a destination? .....	13
Where will your next seniors come from?.....	13
Do you pass the cultural tests? .....	13
Appendix A — Sources and Positioning.....	15
CI&T — AI Doesn't Fix Your SDLC (Grecco, Gaseorowski et al., 2025) .....	15
Angus Gow, Flowfoundry Ltd. — Coordinated AI Development (2026) .....	15
Appendix B — Glossary .....	16

# Executive Summary

---

## *Abundant generation, scarce validation*

The cost of producing software has collapsed. The cost of deciding whether software may move - through review, validation, integration, approval - has not. The industry has largely converged on this diagnosis: the binding constraint in delivery has shifted from execution to decision-making and flow. Two bodies of work now describe the situation well; one strategically, one technically. Neither answers the question that delivery leaders actually face: what does the organization have to become, and how does it get there?

This paper is about that question. It draws on lessons from our own project work - agentic delivery prototyping with regulated and unregulated clients and internal build work - and it is offered as a vision and a hypothesis, not a proven blueprint. Where we have seen work that would conventionally take months completed in days, it was achieved by small, senior, disciplined teams. We treat that caveat not as a footnote but as a finding: the scarce input in agentic delivery is not tooling. It is senior judgment, applied at the right points in the flow, and an organization designed to put it there.

For regulated industries, the question carries an extra edge. In financial services, a decision to advance work is an evidentiary act: it must be possible to demonstrate, months later and to a third party, why something was allowed to move. The constraint is therefore not decision latency but **defensible decision latency**. We call the organizational answer **defensible flow**: delivery in which decisions are explicit, machine-validated where possible, human-gated where necessary, and evidenced always.

The argument runs in four steps. **First**, tools and team-level practices, however good, do not produce the step change; the unit of change is the organization. **Second**, the technical shape of the target is becoming clear: decisions treated as explicit artifacts, validation embedded in the flow, and, eventually, the dissolution of the traditional stage gates around which delivery organizations are built. **Third**, that “eventually” is where transformations live or die: the gates are not just process, they are people, departments, and accountability structures, and redesigning them is the genuinely hard part. **Fourth**, the change must be staged - trust, decision criteria, and orchestration discipline are built in sequence, and skipping stages produces instability, not speed.

Software development is the leading edge of this shift, which is why this paper focuses there. But the same economics, abundant generation, scarce validation, evidence on demand, will reach every form of knowledge work. The organizations that start redesigning now will set the terms for those that follow.

# 1. The Conversation So Far, and the Missing Map

---

Two streams of thinking dominate the current conversation on AI and software delivery, and they operate at different altitudes.

The first is **strategic**. Its core claim: code generation is no longer the constraint; the queues now form at decision points; reviews, approvals, validations, handoffs, and accelerating tasks does not accelerate outcomes when the delay lives between the tasks. Organizations, in this telling, should stop optimizing execution and start orchestrating decisions and flow, progressing through maturity stages from task augmentation through cross-stage coordination toward genuinely agentic delivery. The diagnosis is sound, and we adopt much of it in this paper.

The second stream is **technical and operational**. It describes, often in convincing detail, what disciplined agentic development looks like at team level: explicit context ownership so agents do not collide in shared code; documentation treated as a technical input because agents implement against it at face value; tight coordination rhythms that surface divergence in minutes rather than days; a “harness” of automated controls around every coding agent; commit and review discipline rebuilt for a world where the author may not have written, or fully absorbed, every line. This material is genuinely useful, and Chapter 3 builds on it.

What neither stream provides is the *map between them*. The strategic literature tells leadership why the operating model must change but offers no mechanism and no path. The technical literature offers mechanisms in abundance but stops at the team boundary: it says little about what happens to the departments, governance structures, career paths, and accountability models that surround the team — the things that, in our experience, actually determine whether the technical practices survive contact with the organization.

That missing middle is the subject of this paper. We sketch the technical target briefly, because it is well covered elsewhere and because the reader needs it to follow the argument. We then spend our effort where the literature does not: on the organizational change, including the most consequential and least discussed implication of the technical model — the eventual dissolution of the stage gates around which delivery organizations have been built for forty years — and on a staged hypothesis for how an organization gets from here to there.

A note on standing. This paper is a vision informed by practice, not a report on a completed transformation. The lessons in it come from agentic delivery work we have done with clients and internally - enough to know what the model feels like from the inside, and enough to respect how hard the organizational questions are. We have written it because the map is missing, not because we claim to have finished drawing it.

## 2. The Constraint, Sharpened: Defensible Flow

---

The shared diagnosis can be restated in a paragraph. AI has pushed the marginal cost of producing code, tests, documentation, and prototypes toward zero, while the capacity to validate, integrate, and approve that output has barely moved. Work queues at decision points; lead time is dominated by waiting, not working. And the dynamic is dangerous, not merely inefficient, because of an asymmetry that practitioners of agentic development learn quickly: **misalignment is cheap to create and expensive to unwind**. An agent working from a flawed premise at nine in the morning produces hundreds of lines of internally consistent, plausible, wrong code by ten. The hour it took to create the problem and the day it takes to find and unpick it are wildly disproportionate; and the gap widens with generation speed.

For regulated industries, the diagnosis needs one further turn.

In a regulated institution, a decision to advance work is an **evidentiary act**. When a change moves toward production in a supervised bank or insurer, the decision to let it move carries an obligation that outlasts the release: the institution must be able to demonstrate to an auditor, a risk function, a supervisor what was decided, by whom or by what, against which criteria, on the basis of which validation. Speed that cannot account for itself is not progress; it is liability accumulating with a delay on the invoice.

This reframes the bottleneck. The general literature says it is decision latency. In regulated delivery it is **defensible decision latency**; the time to make a decision and produce the evidence that makes it stand. Today that evidence is mostly reconstructed after the fact: approval emails retrieved, meeting minutes interpreted, review threads excavated. The reconstruction is slow and lossy precisely because the decisions were never captured as decisions; they were ambient, informal, distributed across inboxes and recollections.

**Two consequences follow**, and they frame everything after this point.

**First**, any delivery model that increases the volume and velocity of work while leaving the evidence problem untouched makes a regulated institution *worse off*, however impressive the demonstration. More output through the same implicit decision fabric means longer queues, thinner review attention, and a wider gap between what shipped and what can be accounted for.

**Second**, and this is the opportunity, a delivery model that produces its evidence as a by-product of how work flows removes one of the most expensive and least visible costs in regulated delivery, and converts a compliance burden into a structural advantage. That is what we mean by **defensible flow**: delivery in which decisions are explicit, machine-validated where possible, human-gated where necessary, and evidenced always, not by an army of documenters, but by the design of the flow itself.

The next chapter sketches what such a flow looks like technically. The chapters after it address the harder question of the organization that can run one.

### 3. The Technical Shape of the Target

---

The technical contours of disciplined agentic delivery are emerging clearly across the industry, and this chapter synthesizes them briefly, as orientation, not as novelty. The organizing idea is simple to state: **decisions become first-class artifacts**. Instead of being ambient events; things that happen in meetings, in reviewers' heads, in approval emails, the decisions that advance work are specified before work begins, evaluated continuously while it runs, gated where judgment is required, and recorded durably when work moves.

In practice, the pipeline has six elements.

**Intent and specification.** Work begins with a structured specification: the outcome required, the constraints that bind, the acceptance and decision criteria that will be applied. The specification is the contract that agents execute and validators judge against. Writing it precisely is skilled work, a point we return to, because it turns out to be *the* skilled work.

**An explicit plan.** The specification is decomposed into a machine-readable plan: a graph of work units with dependencies, each carrying its own definition of done. The plan is reviewed before execution; a cheap point to catch expensive errors, given the asymmetry of Chapter 2, and persists as the reference against which divergence is detected as divergence rather than discovered as surprise.

**Orchestrated execution under guardrails.** Agents execute work units, in parallel where dependencies allow, inside explicit boundaries: what they may touch, which interfaces they may not change without escalation, which standards apply. Guardrails convert “implement this feature” from an open-ended act into a bounded one whose blast radius is known in advance. Human navigators supervise, and the quality of that supervision is, in our experience, the single strongest determinant of output quality.

**Continuous validation, in two layers.** Validation runs with execution, not after it. *Computational controls*, deterministic, fast, cheap: type checks, linting, structural and architectural boundary tests, coverage gates, run on every change, inside the agent's working loop. *Inferential controls*, judgment-based: review agents, semantic analysis, AI-as-judge evaluation against the specification's intent, run selectively, where stakes justify cost. Every validation event is recorded against the work it validated: which checks, against which criteria, when, with what result.

**Sharp human gates where judgment is required.** Where a decision genuinely requires human judgment, and in regulated delivery, some always will, the flow presents a structured gate: the work, its validation history, its divergences from plan,

and the specific question being asked. This inverts today's review burden. Instead of asking a senior person to reconstruct context from a four-hundred-line diff, it asks a precise question against assembled evidence. The decision, the rationale, and the decider are captured as part of the record.

**A durable evidentiary record.** Everything above accumulates into a single artifact per unit of delivered change, call it a **case file**: originating intent, specification, plan and revisions, execution trace, every validation event, every exception and its resolution, every human gate decision. This is "defensible" operationalized. It is what an auditor, a risk reviewer, or a future engineer consults; not a reconstruction, but the record as it was made. It is also the institutional memory that survives staff turnover, and the context the next agent loads before touching the same component.

Three properties of this model deserve emphasis, because they carry the organizational argument of the next chapter.

**First: the gates dissolve, eventually.** Notice what the model does to traditional stage-gate governance. The checkpoints that structure conventional delivery; QA sign-off, architecture review board, change advisory board, compliance review, release approval, exist because validation was expensive, manual, and episodic, so it was batched at gates. When validation is continuous, embedded, and recorded, the function of most gates is performed better inside the flow than at its boundaries. What remains is a small number of sharp human decision points; far fewer, far better informed. The eventual end state is not "faster gates"; it is, for most decision classes, *no gates*, replaced by continuously evaluated confidence and a durable record. The word *eventually* is doing essential work in that sentence, and Chapter 5 explains why. But the direction of travel should be stated plainly, because the organizational consequences in Chapter 4 follow from it.

**Second: a hard ceiling exists, and honesty about it matters.** There is an unsolved problem across the industry: tests generated by the same agent that wrote the code tend to verify the agent's interpretation of the requirement, not the requirement itself. For correctness-critical logic; money movement, regulatory calculation, access control, cryptography, agent-generated validation is insufficient by construction. Such logic carries mandatory human-authored specification tests and a mandatory human gate, indefinitely or until genuinely independent validation exists. Any model, or vendor, that claims past this ceiling should be treated with suspicion.

**Third: delegation is a governed map, not a dial.** "How autonomous?" is the wrong question. The right one is "which decision classes sit where?"; *delegated* (machine decides, record kept), *assisted* (machine recommends with evidence, human decides), or *reserved* (human decides, machine assembles the file). The placement of each decision class is itself a governed, recorded choice that moves only as evidence accumulates. In a regulated institution, this human gate map is the document the risk

function will actually want to see, and producing it deliberately is far better than having it inferred from incident reports.

All of this is buildable. Much of it, in various forms, is being built; by us and by others. The technology is not the hard part. The next chapter is about the hard part.

## 4. The Organizational Change: the Hard Part

---

Here is the uncomfortable arithmetic of this transformation: the technical layer of Chapter 3 can be stood up in months. The organization that can run it takes years. That asymmetry; not model capability, not tooling maturity, is where transformations stall, and it is the least examined part of the current conversation. What follows is our hypothesis for what actually has to change, stated as concretely as we can. We hold it as a hypothesis informed by project experience, not as settled knowledge; the honest state of the industry is that nobody has completed this transformation at enterprise scale.

### 4.1 The gates are people

Start with the implication Chapter 3 stated and most discussions avoid: traditional stage gates dissolve, eventually, into continuous validation plus a small number of sharp human decision points.

Stage gates are not just process boxes on a delivery diagram. They are departments, roles, careers, and accountability structures. QA functions exist to operate the testing gate. Release management exists to operate the deployment gate. Change advisory boards exist to operate the production-change gate. Architecture review boards, security sign-off, compliance review; each is an organizational answer to the same historical fact: validation was expensive and episodic, so it was batched at checkpoints, and people were employed to run the checkpoints.

When validation becomes continuous and embedded, the work of these functions does not disappear, it transforms, and the transformation is profound. The QA function's future is not executing tests at a gate; it is **designing and assuring the validation system itself**, defining what the computational and inferential controls must check, hunting the edge cases automation misses, owning the integrity of the mechanism that replaced the gate. The change-management function's future is not convening boards; it is **owning the human gate map**, the governed classification of which decisions are delegated, assisted, and reserved, and the evidence standards that allow a decision class to move bands. The compliance and security review functions move from inspecting outputs at checkpoints to **encoding their criteria as machine-enforceable controls** inside the flow, and auditing the controls rather than the artifacts.

In every case the pattern is the same: *from operating the gate to engineering the criteria*. It is a move up the value chain and it is experienced as a threat before it is experienced as a promotion. People whose professional identity and organizational power derive from being the checkpoint do not cheer when the checkpoint dissolves. A transformation plan that does not name these functions, involve them early in designing the controls that replace their gates, and give them an explicit and respected destination role will be quietly strangled by the very people it needs most. In regulated institutions, where the second line of defence holds real authority, this is doubly true: the risk function must be a co-designer of the embedded control model, not its last surprised reviewer.

## 4.2 The seniority inversion

The second structural change concerns who does what, and it runs opposite to the intuition that AI primarily threatens senior, expensive people.

In our project work, the consistent lesson has been this: where agentic delivery produced dramatic compression; work conventionally measured in months completed in days, it was in the hands of small, senior, disciplined teams. People who could write specifications precise enough to act on, review plans before execution and catch the expensive error early, supervise agent output with genuinely critical attention across the full stack, and stop the moment something diverged rather than letting plausible-looking output accumulate. The method amplified their judgment. We have seen no evidence that the method alone, handed to a team without that judgment and discipline, produces anything similar, and the well-documented adoption dip most teams experience suggests it does not.

The organizational consequence is an inversion of the classical delivery pyramid. For decades, leverage came from junior volume: many hands executing well-specified tasks under a thin senior layer. In agentic delivery, routine execution is what the agents absorb. What becomes scarce and therefore central is the senior skill set the industry once called the analyst-programme: deep domain understanding, specification precision, critical evaluation of work one did not write, and a live architectural model of where the system is going. The supervising role - the **navigator** - is a senior role, not a junior one. The developer who treats agentic tools as an opportunity to disengage from the code is the developer whose sessions produce the most expensive problems.

This creates the most awkward question in the whole transformation: **where do the next seniors come from?** The traditional path, juniors learning by writing volumes of routine code, runs straight through the territory the agents now occupy. An organization that simply stops hiring juniors is eating its seed corn; one that hires them into the old path is training them for work that is disappearing. The development path has to be deliberately rebuilt: juniors learning by reviewing agent output against specifications, by writing and refining specifications under supervision, by working

inside the validation and harness layer where the system's failure modes are most visible, and by acquiring domain depth far earlier than the old path required. No one has fully solved this. Organizations that start experimenting with it now will hold a structural talent advantage within five years over those that wait.

### **4.3 The coordination fabric becomes an organizational design problem**

The technical literature describes, correctly, the team-level disciplines that agentic velocity demands: explicit ownership of code territories so parallel agent sessions do not collide; continuous, low-overhead broadcast of intent and state changes, because divergence at AI speed must be caught in minutes, not at tomorrow's standup; documentation treated as a technical input with owners and update obligations, because agents implement against it at face value and stale documents now produce confidently wrong code at volume; lightweight ceremonies, session context loads, drift reviews, architecture checkpoints, that keep implementation and intent aligned.

What is less discussed is that, beyond a handful of teams, none of this survives as voluntary craft. Conventions that one excellent team maintains by culture must, at organizational scale, become **designed infrastructure with ownership**: the context and documentation layer becomes a maintained asset with a named owner, not housekeeping; the harness, controls, guardrails, templates, the case-file machinery, becomes a platform product with a roadmap and a team, not per-project scaffolding rebuilt and abandoned; coordination norms become onboarding curriculum and management expectation, not folklore. The platform-engineering function, already ascendant, becomes the natural home: a small team with deep expertise owning the golden paths and the validation fabric, enabling a larger population of full-stack practitioners to work safely within them. Organizations that leave the harness as a per-team hobby will watch their early wins fragment into islands; the harness is where learning compounds, and compounding requires ownership.

### **4.4 Measures and incentives move, or nothing else does**

Delivery organizations steer by what they measure, and the inherited measures, velocity, utilization, on-time gate passage, are all artifacts of the era when execution was scarce. Each one, kept in place, actively fights the new model: velocity rewards volume of output precisely when output has become cheap and absorbing it has become the constraint; utilization rewards keeping people busy precisely when the system's health depends on validation capacity and slack at decision points; gate-passage metrics reward the ritual the model is dismantling.

The replacement measures follow from the argument: **flow efficiency** (the share of lead time spent on active work rather than waiting), **decision latency by decision class, rework rate** (the misalignment asymmetry made visible and priced), and, the measure specific to defensible flow or **time-to-evidence**, how quickly the

organization can produce the complete, accurate record for any delivered change. In the target state the answer is “immediately, it already exists,” and the distance from that answer is a direct measure of transformation progress. Incentives, budgets, and status conversations must move to these measures early, not as the final polish but as one of the first interventions, because middle management optimizes what the dashboard shows, and middle management is where this transformation is won or lost.

#### **4.5 The cultural preconditions are testable**

Finally, the soft layer that decides everything - which can be made concrete with three tests. Can a developer in your organization say “I don't understand this code” without career damage? Because honest attestation about AI-generated output is the foundation of the entire review model, and in cultures where admitting confusion is penalized, attestation will be quietly fictionalized. Does a divergence alert get attention within minutes, with the unremarkable, blame-free quality of a routine correction? Because the economics of the asymmetry depend on correction speed, and correction speed depends on it being socially costless to raise a hand. And is “the agent found a better path than the plan” treated as a welcome, recordable event rather than a violation? Because a system of explicit plans only stays honest if good deviation is captured rather than hidden. Organizations that fail these tests have culture work to do before tooling work; the tooling will only make the culture's failure modes faster.

## **5. Getting There: a Staged Hypothesis**

---

If Chapter 4 is right about what must change, the path matters as much as the destination. The single most reliable finding across early adopters is that **stages cannot be skipped**. Trust, decision criteria, and orchestration discipline are built sequentially; autonomy granted before the evidence exists produces instability and incident-driven retrenchment, which is slower than never having hurried. What follows is our hypothesis for the sequence, organized around a principle each stage must honor: *gates are dismantled only after the thing that replaces them has earned its authority*.

### **Stage 1: Build the fabric before scaling the agents**

The counterintuitive first move is that the earliest investment is not in agent adoption but in what surrounds it. A small senior group, harness builders, establishes the foundations: the context and documentation layer agents will consume, the computational controls that will run on every change, specification and plan templates, the coordination disciplines, the initial human gate map with every consequential decision class still in the reserved band. Broader agent use begins only on this fabric, and deliberately on low-stakes work first; tests, documentation, well-understood components, while teams climb the learning curve. Leadership's job in this stage is

expectation-setting: the productivity dip is real, typically weeks for well-prepared teams, and announcing it in advance is the difference between a learning curve and a credibility crisis. *Exit evidence*: controls run on every change; the case-file record exists end-to-end for real work; at least one team beyond the founding group has reached working discipline.

## **Stage 2: Run the new model inside the old governance**

This is the stage most plans skip, and the reason most transformations stall. The embedded validation flow runs in parallel with the existing gates; the QA sign-offs, the boards, the reviews all still happen, while the organization accumulates the comparative evidence: what did the continuous controls catch, what did the gate catch that the controls missed, where do they disagree. This feels like duplication because it is duplication; temporarily and on purpose: it is how the replacement earns its authority, and it is how the gate-operating functions participate in building, and learn to trust, and begin to own the thing that will transform their roles. In regulated institutions this stage has a second purpose: it produces the evidence file that the risk function and, ultimately, the supervisor will require before any gate is retired. The role transformations of section 4.1 are negotiated and begun here, function by function, with explicit destination roles. *Exit evidence*: for specific decision classes, the embedded controls demonstrably match or outperform the gate they shadow, over a defined period, with the gap documented.

## **Stage 3: Retire gates by decision class, by evidence**

Gates are dismantled the way they were validated: one decision class at a time, moving bands on the human gate map only on accumulated evidence, with reversal as an explicit, non-punitive option. Low-risk classes move first; correctness-critical classes, the behaviour-validation ceiling of Chapter 3, may remain reserved indefinitely, and saying so plainly is a strength, not a concession. In parallel, the measures of section 4.4 replace the legacy dashboard, the platform function takes permanent ownership of the harness, and the rebuilt junior development path takes its first cohort. The end state is not a gateless free-for-all; it is a delivery organization with continuous, evidenced confidence and a short list of sharp, well-posed human decisions; defensible flow as the normal way work moves.

Three failure modes recur often enough to name. **Tool-first adoption**, buying agents for everyone before the fabric exists, produces local wins, systemic queues, and a backlash that poisons the second attempt. **Gate removal by decree**, declaring the boards abolished before the embedded controls have earned authority, produces either an incident that triggers wholesale retrenchment, or shadow gates that reappear informally with none of the old transparency. And **transformation by pilot**, a permanently celebrated island that never converts its insights into operating-model change, produces conference talks. The staged path is slower to start than any of the three and faster to arrive than all of them.

## 6. What Leaders Should Do Now

---

For a CIO, CTO, COO, or head of delivery reading this, the honest summary is: the technology will keep improving without your help; the organization will not. The leverage is in starting the organizational work early, because it has the longest lead time. Five questions are worth putting to your own organization this quarter.

### **Where are your decisions, actually?**

Map the decision points in your delivery flow; formal gates and informal ones. And for each, ask what it validates, what evidence it produces, and how long work waits in front of it. Most organizations have never drawn this map and are surprised by what it shows. It is also the baseline against which any transformation must be measured.

### **Could you produce the file?**

Pick a recent material change in production and ask for the complete record: who decided what, against which criteria, on what validation. Time how long the reconstruction takes. That duration is your current time-to-evidence, and the gap between it and “immediately” is the size of the defensible-flow opportunity in your institution.

### **Do your gate-operating functions have a destination?**

Ask whether your QA, change-management, release, and review functions have heard a credible account of their future role, criteria engineering, gate-map ownership, validation assurance, or whether they are quietly assuming the transformation is something being done to them. Their early involvement is not change-management courtesy; it is a structural dependency.

### **Where will your next seniors come from?**

If routine execution is what the agents absorb, examine your junior development path and ask what, concretely, replaces learning-by-volume. If the answer is silence, that silence has a five-year fuse.

### **Do you pass the cultural tests?**

Can “I don't understand this code” be said safely? Does a raised hand get a five-minute response or a retrospective? The answers predict your adoption curve better than any tooling evaluation.

None of this requires waiting for the technology to settle, and all of it pays off regardless of which tools win, because decision maps, evidence discipline, role clarity, talent pipelines, and culture are tool-agnostic assets.

A closing observation. This paper has focused on software development because software is where the new economics, abundant generation, scarce validation, evidence on demand, arrived first and where the organizational responses are furthest

along. But nothing in the argument is specific to code. Every discipline whose product is structured knowledge work, analysis, advisory, legal drafting, financial reporting, consulting itself, will meet the same inversion, on a short delay. The organizations that learn, in software, how to dissolve gates into evidenced flow and how to redesign roles around scarce judgment will not just deliver software better. They will hold the template for what comes next.

## Appendix: Sources and Positioning

---

This paper builds deliberately on two published works, and the intellectual debts should be explicit.

### **CI&T: AI Doesn't Fix Your SDLC (Grecco, Gaseorowski et al., 2025)**

We adopt its central diagnosis; that the constraint has moved from code production to decision-making and flow, and that AI applied to execution alone amplifies queues rather than outcomes, and its staged-maturity logic, including the warning that skipping stages produces instability. We depart from it in three ways: we supply a concrete mechanism (decisions as artifacts, embedded two-layer validation, the case file) where it remains conceptual; we center the organizational change, gate dissolution, role transformation, talent pipeline, measures, which it treats only in passing; and we add the regulated dimension, where the constraint is not decision latency but defensible decision latency. We do not adopt its quantitative multipliers, which are presented there without supporting evidence.

### **Angus Gow, Flowfoundry Ltd: Coordinated AI Development (2026)**

We adopt its account of the misalignment asymmetry; its team-level coordination model (context ownership, continuous low-overhead communication, drift discipline); its harness framing, feedforward guides and feedback sensors, computational versus inferential controls, which we use directly in Chapter 3; its honesty about the unsolved behaviour-validation problem; and its readiness and adoption-curve analysis. We depart from it by elevating these from team practice to organizational design, arguing that beyond a few teams, the disciplines it describes survive only as owned infrastructure, and by drawing out the implication it leaves implicit: that continuous embedded validation eventually dissolves the stage-gate structure of the delivery organization, with all the consequences of Chapter 4.

The synthesis, the defensible-flow framing, the organizational-change hypothesis, and the staged path are this paper's own, informed by the authors' project work in agentic delivery with financial-services clients and internally.

## Appendix B: Glossary

---

**Defensible flow.** Delivery in which decisions are explicit, machine-validated where possible, human-gated where necessary, and evidenced always.

**Defensible decision latency.** The time to make a decision and produce the evidence that makes it stand.

**Specification.** The structured statement of outcome, constraints, and acceptance criteria against which execution and validation are judged.

**Plan.** The machine-readable decomposition of a specification into dependent work units, each with its own definition of done.

**Guardrail.** An explicit boundary on agent action (scope, interfaces, standards) that bounds blast radius in advance.

**Computational control.** Deterministic, cheap validation run continuously (types, lint, structural and boundary tests, coverage).

**Inferential control.** Judgment-based validation (review agents, AI-as-judge) applied selectively where stakes justify cost.

**Human gate.** A structured decision point presenting assembled evidence and a precise question to an accountable person.

**Case file.** The durable per-change record binding intent, specification, plan, execution trace, validations, exceptions, and gate decisions.

**Human gate map.** The governed classification of decision classes into delegated, assisted, and reserved bands, moved only on accumulated evidence.

**Navigator.** The senior practitioner role supervising agentic execution with critical attention.

**Harness.** The system of controls, guardrails, templates, and context infrastructure surrounding agentic execution; at organizational scale, a platform product with an owner.

**Behaviour-validation ceiling.** The unsolved problem that validation generated by the producing agent verifies the agent's interpretation rather than the requirement; caps delegation for correctness-critical logic.

**Time-to-evidence.** How quickly the organization can produce the complete record for any delivered change; in the target state, immediate.