Four-Level Model of Organisational Agility

An Application of the Adaptive Fitness System

By Mario Aiello

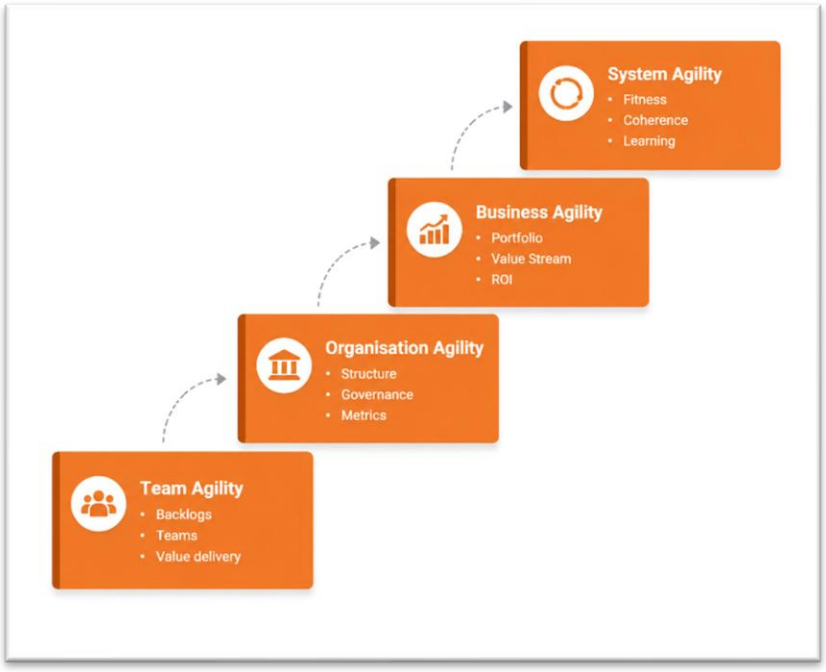

Abstract

This paper presents the Four-Level Model of Organisational Agility (4LM), a diagnostic and developmental approach grounded in the Adaptive Fitness System (AFS) and the Fitness-Validation-Intelligence Lens (FVI). The model identifies four levels at which organisational agility must be developed, tested for coherence, and sustained: Team Agility, Organisational Agility, Business Agility, and System Agility. Each level is defined by three diagnostic pillars, each aligned to a dimension of AFS Fitness. The paper argues that sustainable agility is not achieved by optimising individual levels independently, but through coherent fitness across all four — and that System Agility, the fourth level, is not a sequential step but an emergent, continuous condition whose presence or absence manifests across the entire model.

The approach is applicable across organisational contexts. Software development organisations are used throughout as the primary illustrative domain. The paper supersedes three predecessor documents — the Three-Level Model Grounded in AFS-FVI, How the 3LM Fits the AFS Ecosystem, and the Three-Level Model Practitioner Guide — consolidating their content into a single canonical reference.

1. Introduction

Over two decades of practice in organisational transformation, a consistent pattern emerges: organisations invest heavily in agility at one level while ignoring the conditions that would allow agility

to function at another. Teams are taught to iterate; architectures are left monolithic. Governance is restructured; strategy remains opaque. Metrics are refined; learning never reaches the parts of the organisation that need it.

The diagnosis is not that agile methods fail. It is that agility requires coherence across levels that most transformation approaches treat as separate problems. A team cannot be agile if the organisation prevents it from acting autonomously. An organisation cannot be agile if business strategy is misaligned or unresponsive. And neither team nor organisational agility compounds into durable capability if the system as a whole lacks the meta-capacity to assess its own fitness, maintain coherence under pressure, and learn at scale.

The Four-Level Model of Organisational Agility (4LM) addresses this directly. It provides a structured diagnostic approach for identifying where coherence breaks down — and a developmental framework for restoring it. The model is grounded in two foundational ideas from the Adaptive Fitness System.

First, the AFS Fitness Sequence: change efforts must be fit for purpose, fit for context, fit for execution, and fit for improvement — in that order. Skipping steps creates interventions that are locally coherent but systemically inert. Second, the Fitness-Validation-Intelligence (FVI) Lens: fitness must be validated through reverse reasoning from desired outcomes, and sustained through adaptive intelligence — a continuous sense–learn–adapt loop that keeps the organisation coherent as conditions shift.

The 4LM operationalises these ideas across four levels. Each level is defined by three pillars, each pillar mapped to a dimension of the AFS Fitness Sequence. The model does not prescribe solutions. It surfaces the actual constraint holding agility back, and provides the conceptual structure to address it coherently.

> A note on scope
>
> *This paper uses software development organisations as its primary illustrative domain. The diagnostic language — sprints, CI/CD, deployment frequency, value streams — is drawn from that context. The underlying model is not software-specific. The same four levels, three pillars, and AFS alignment apply wherever organisational agility is being developed: financial services, healthcare, public sector, product organisations, or any context where coherent adaptation to changing conditions is a strategic requirement.*

2. Theoretical Foundation

2.1 The Adaptive Fitness System

The Adaptive Fitness System (AFS) is an approach to organisational change developed by Mario Aiello through 25 years of practice in transformation, primarily in European banking and financial services. It reframes the central question of organisational development from "which method should we adopt?" to "what are we actually fit for, and how do we know?"

The AFS is built on a Fitness Sequence of four dimensions, applied in order:

- Fit for Purpose — Does the proposed change address the actual problem? Is the problem well-understood, or assumed?
- Fit for Context — Does the approach account for the specific conditions in which it will operate — structure, culture, constraints, history?

- Fit for Execution — Can the organisation actually implement this? Do the capabilities, resources, and conditions for execution exist?
- Fit for Improvement — Is there a feedback mechanism that will allow the approach to evolve as conditions change?

These are not a checklist. They are a sequenced discipline. An approach that is fit for purpose but not for context will be correctly motivated but misdirected. One that is fit for context and execution but not improvement will produce results that degrade as conditions shift. The sequence matters because skipping a step does not eliminate the problem it was designed to surface — it merely defers it.

2.2 The Fitness-Validation-Intelligence Lens

The FVI Lens is the analytical layer of AFS. It adds three practices to the Fitness Sequence:

- Fitness — applying the four-dimension sequence before committing to any intervention
- Reverse Validation — working backward from a desired outcome to identify what must be true at each prior level for that outcome to be achievable; this surfaces hidden dependencies and exposes weak assumptions before resources are committed
- Adaptive Intelligence — maintaining a continuous sense–learn–adapt loop throughout implementation, so that the approach evolves in response to what actually happens rather than what was planned

Reverse Validation is particularly important to the 4LM. The model's four levels are sequentially dependent: Team Agility is a precondition for Organisational Agility, which is a precondition for Business Agility. Reverse Validation forces the question in the other direction — if Business Agility is the desired outcome, what must be true at the organisational level? If organisational coherence is required, what must be true at the team level? This exposes the actual constraint rather than the most visible symptom.

2.3 The 4LM as AFS Application

The 4LM is a domain-specific operationalisation of AFS. AFS provides the meta-layer — the compass — asking whether change efforts are fit for purpose, context, execution, and improvement. The 4LM provides the map: in an organisational agility context, this is what those fitness dimensions look like at each level, what to diagnose, what failure modes to expect, and how the levels depend on each other.

The relationship is not supplementary. The 4LM without AFS is a maturity model — a set of levels to climb. The 4LM grounded in AFS is a fitness model — a structured approach to diagnosing what is actually preventing coherent agility and addressing it in sequence. The difference is consequential: maturity models tend to drive adoption of practices; fitness models drive understanding of conditions.

3. Model Architecture

The 4LM comprises four levels, each representing a distinct organisational domain in which agility must be developed. Each level contains three diagnostic pillars, each aligned to a dimension of AFS Fitness. The levels are sequentially dependent in the developmental direction: Team Agility is the foundation; System Agility is the emergent condition.

> Sequential dependency vs. continuous condition
>
> *The first three levels — Team, Organisational, and Business Agility — are sequentially dependent. A team cannot exercise genuine agility if the organisation prevents autonomous action. An organisation cannot be coherently agile if business strategy is opaque or unresponsive. These dependencies run upward: each level requires the level below it to be functional.*
>
> *System Agility is different. It is not a fourth step in a sequence. It is a continuous condition — an emergent property — that either manifests across the whole model or does not. Its three pillars (Meta-Fitness, Cross-Level Coherence, and Adaptive Learning) are not achieved once and maintained; they are either present as active capacities or absent as structural gaps, regardless of how well any individual level is functioning.*

Table 1: 4LM Architecture Overview

Level	Defining Question	Three Pillars	AFS Alignment
Team Agility	Can individual teams iterate, learn, and deliver value consistently?	Clarity of Intent / Execution Capability / Delivery Commitment	Fit for Purpose / Fit for Execution / Fit for Improvement
Organisational Agility	Can the organisation enable teams to act coherently at scale?	Coherent Structure / Suitable Governance / Meaningful Metrics	Fit for Context / Fit for Purpose / Fit for Improvement
Business Agility	Can the business sense, learn, and respond to strategic opportunity?	Portfolio Management / Effective Value Streams / Clear ROI	Fit for Purpose / Fit for Execution / Fit for Improvement
System Agility	Does the system maintain meta-fitness, coherence, and learning across all levels?	Meta-Fitness / Cross-Level Coherence / Adaptive Learning	AFS Fitness (meta) / FVI Coherence / Adaptive Intelligence

4. Team Agility

A team is agile when it understands what it is building and why, can execute that work with consistent quality, and commits to a sustainable rhythm of delivery and learning. These are not independent virtues; they are sequentially dependent conditions. Execution capability without clarity of intent produces high-quality solutions to the wrong problem. Delivery commitment without execution capability produces a reliable cadence toward an accumulation of technical debt.

4.1 Clarity of Intent — Fit for Purpose

The first condition for team agility is that the team can articulate what it is building, why it matters, and how success will be recognised. This is not a communication requirement; it is a fitness requirement. Work that is not fit for purpose — not grounded in a clear understanding of the problem being solved — will generate rework regardless of how well it is executed.

Diagnostic questions

- Can the team explain the user or business problem being solved, not just the feature?
- Are acceptance criteria defined before work begins, or discovered during it?

- Does the team understand how its work connects to organisational goals?
- Is the backlog prioritised against value outcomes, or against urgency and noise?
- Can the team decline work that does not fit the stated purpose?

Failure modes

Backlog ownership concentrated in a single person with no team input. Priorities that shift within a sprint cycle without strategic justification. Stories written as technical specifications without business context. No feedback mechanism from users or stakeholders back to the team's work queue.

Observable signal of health

The team can say no to work that does not fit the purpose. Sprint planning conversations centre on value delivered, not tasks completed. Stakeholder and team definitions of done converge rather than conflict.

4.2 Execution Capability — Fit for Execution

A team may understand clearly what it is building and still be unable to deliver it reliably. Execution capability is the structural condition that converts intent into delivery. It encompasses technical skill, tooling, architecture, and the organisational permissions that allow the team to act.

Diagnostic questions

- Can the team deploy without manual approval or extensive coordination with other teams?
- Is testing automated at multiple levels — unit, integration, system?
- Is technical knowledge distributed across the team, or concentrated in specialists?
- Can the team identify and address technical impediments without external intervention?
- Is technical debt being managed incrementally, or accumulating?

Failure modes

Monolithic architecture requiring all teams to coordinate deployments. Manual testing creating release bottlenecks. Knowledge silos generating single points of failure. Absence of CI/CD infrastructure. Technical debt treated as a backlog item perpetually deferred.

Observable signal of health

Deployment is routine rather than a crisis event. New team members contribute meaningfully within weeks rather than months. Technical problems are surfaced and addressed rather than concealed.

4.3 Delivery Commitment — Fit for Improvement

Delivery commitment is the capacity to sustain a reliable rhythm of work and learning over time. It is not about velocity. It is about cadence, feedback, and the psychological safety required for honest reflection. A team that cannot commit to a sustainable rhythm cannot learn from it.

Diagnostic questions

- Does the team work in fixed-length iterations with a predictable review cadence?
- Are feedback loops active — do users and stakeholders regularly interact with working software?
- Do retrospectives produce changes that are actually implemented?
- Is the pace sustainable, or are individuals consistently operating beyond capacity?

- Are failures treated as sources of learning or as grounds for blame?

Failure modes

Constant interruption breaking sprint rhythm. Retrospectives treated as compliance exercises. Feedback loops measured in weeks rather than days. Top-down process controls that prevent the team from adapting how it works. Blame culture suppressing honest problem identification.

Observable signal of health

Team morale is stable. Velocity is predictable. People experiment and report problems early. Retrospective action items are completed rather than carried indefinitely.

Table 2: Team Agility Diagnostic Summary

Dimension	Fitness Alignment	Key Diagnostic Question	Failure Mode
Clarity of Intent	Fit for Purpose	Can the team explain the problem, not just the feature?	Rework, context-switching, misalignment with stakeholders
Execution Capability	Fit for Execution	Can the team deploy without coordination or manual approval?	Bottlenecks, knowledge silos, accumulating technical debt
Delivery Commitment	Fit for Improvement	Is the pace sustainable with learning built in?	Burnout, stagnation, blame culture, retrospective theater

5. Organisational Agility

An organisation is agile when it enables multiple teams to act coherently, with minimal forced coordination, clear decision authority, and feedback mechanisms that connect execution to strategy. Team-level agility is necessary but insufficient. A high-functioning team operating inside an incoherent organisation will exhaust its energy managing dependencies that should not exist.

5.1 Coherent Structure — Fit for Context

Structure is not neutral. Organisational and technical structures either enable autonomous action or prevent it. Most organisations have structures that made sense at a previous scale or strategic context and have not been revisited. The result is coordination overhead that is treated as inevitable rather than as a structural choice.

Diagnostic questions

- Are teams organised around value streams or business capabilities, rather than technical functions?
- Can a single team deliver end-to-end value, or is value systematically split across teams?
- Is the technical architecture aligned with team boundaries — loosely coupled, with clear interfaces?
- Are cross-team dependencies intentional and minimised, or accidental and pervasive?
- Do teams have clear ownership of the components and services they operate?

Failure modes

Functional silos — Frontend, Backend, QA, Operations as separate teams — that force handoffs for every delivery. Monolithic architectures that require all teams to coordinate any deployment. Shared services owned by bottleneck teams that have no structural incentive to prioritise other teams' needs. Unclear ownership generating constant negotiation about who is responsible for what.

Observable signal of health

Teams can deploy independently. Most features can be completed within one or two teams. Architecture decisions are made at the team level within guardrails, not escalated to central committees.

5.2 Suitable Governance — Fit for Purpose

Governance is necessary. The question is not whether to govern but what to govern, at what level, and by what logic. Most organisational governance has accumulated reactively — rules added after incidents, approval layers added after failures — without being reviewed for whether the overhead it creates is proportionate to the risk it mitigates.

Diagnostic questions

- Who makes decisions, and at what level? Are most decisions made locally?
- Are governance rules grounded in assessed risk, or in institutional fear?
- How long do approvals take? Hours, days, or weeks?
- Are governance rules known and understood by the teams they govern?
- When conflicts arise between teams, can they be resolved without escalation?

Failure modes

Approval processes requiring multiple sequential hand-offs. Governance rules disconnected from actual risk exposure. Decision authority ambiguous enough that escalation becomes the default. Rules that are rarely enforced but create process theater when they are. Governance overhead that consumes more organisational capacity than the risk it prevents.

Observable signal of health

Teams know the rules and understand the reasoning behind them. Escalations are rare and resolved quickly. Governance is reviewed and updated rather than accumulated.

5.3 Meaningful Metrics — Fit for Improvement

Metrics are an organisational feedback mechanism. Their quality determines the quality of the learning they enable. Most organisations measure activity — throughput, velocity, tickets closed — because activity is easy to observe. Outcome metrics — value delivered, customer impact, strategic progress — require more effort to define and instrument, and more tolerance for findings that challenge current priorities.

Diagnostic questions

- What metrics does the organisation actually act on — throughput, or outcomes?
- Do teams understand their metrics and see how their work connects to them?
- Are metrics reviewed with sufficient frequency to inform decisions rather than just report history?

- When metrics degrade, does the organisation investigate systemically or respond with blame?
- Is there a feedback loop from metrics back to strategic priorities?

Failure modes

Metrics focused on activity rather than outcomes. Dashboards visible only to leadership, not to teams. Metrics that change frequently, preventing trend analysis. Metrics used to identify individuals to blame rather than systems to improve. Measurement lag so long that the data describes conditions that no longer exist.

Observable signal of health

Teams make decisions based on data. Experiments are run, results are analysed, and learnings are shared across the organisation rather than retained within teams.

Table 3: Organisational Agility Diagnostic Summary

Dimension	Fitness Alignment	Key Diagnostic Question	Failure Mode
Coherent Structure	Fit for Context	Can teams act autonomously, or does structure force coordination?	Dependency overhead, deployment bottlenecks, negotiation theater
Suitable Governance	Fit for Purpose	Are decisions made locally and quickly, with risk-proportionate oversight?	Approval bottlenecks, escalation loops, governance as control theater
Meaningful Metrics	Fit for Improvement	Are outcomes measured and acted upon to drive organisational learning?	Activity theater, no strategic insight, metric-driven blame

6. Business Agility

Business agility is the capacity to sense market opportunity, align resources coherently, and iterate strategy based on what execution reveals. It is the level at which agility connects to competitive relevance. An organisation can have highly capable teams and coherent structure and still fail at business agility if strategy is opaque, prioritisation is driven by noise, and learning from execution never reaches the people who set direction.

6.1 Portfolio Management — Fit for Purpose

Portfolio management answers the question that every team eventually asks: why are we building this instead of something else? Without a coherent answer to that question, prioritisation defaults to urgency, noise, or organisational power. The result is fragmented resources, competing priorities, and the systematic underinvestment in work that matters strategically but lacks a vocal internal advocate.

Diagnostic questions

- Is there a clear strategy that teams can articulate and connect to their work?
- Is work prioritised against strategic criteria and expected ROI, or against urgency and demand?
- How often does strategy change? Never suggests brittleness; constantly suggests incoherence.

- Are trade-offs made explicit — this work instead of that work?
- Is deprioritisation as visible as prioritisation?

Failure modes

Strategy unclear or not communicated below leadership level. Prioritisation driven by the most vocal stakeholder. No ROI evaluation before work is committed. Portfolio decisions made annually and treated as fixed, regardless of market signals. Resource allocation spread thin across too many simultaneous initiatives.

Observable signal of health

New requests are evaluated against strategic criteria before being committed. Teams understand why they are building what they are building. Portfolio decisions are made in days, not months.

6.2 Effective Value Streams — Fit for Execution

A value stream is the path from strategic intent to realised user value. Defining and managing value streams is the organisational act of making strategy executable. Where value streams are poorly defined, strategy fragments into a collection of disconnected team activities. Where they are well-defined, strategic intent flows through the organisation with minimal translation loss.

Diagnostic questions

- Are value streams clearly defined and named?
- Is there a single owner accountable for each stream's outcomes?
- How many teams are involved in delivering a single value stream, and are cross-stream dependencies explicit?
- Where do delays occur within the stream, and are they visible or hidden by local team metrics?
- Can value streams adapt to changes in scope or strategy without extensive renegotiation?

Failure modes

Value streams poorly defined or with unclear ownership. Handoffs between functional units creating systematic delays. Dependencies on shared services owned by teams with no accountability for stream-level outcomes. End-to-end flow invisible — each team appears fast, but the stream is slow. Scope changes mid-stream generating cascading rework.

Observable signal of health

Strategic initiatives flow end-to-end with minimal coordination overhead. Cycle time from idea to user impact is measured and improving. Value stream retrospectives identify systemic improvements across team boundaries.

6.3 Clear ROI — Fit for Improvement

ROI in this context is not a financial calculation; it is an epistemic commitment. It is the organisation's willingness to define what success looks like before committing resources, measure whether it was achieved after execution, and feed that learning back into strategy. Without it, organisations can be highly active and strategically blind simultaneously.

Diagnostic questions

- Do initiatives have defined success criteria before launch, not retrospectively?

- Is post-launch measurement happening — usage, adoption, satisfaction, revenue impact?
- When results are disappointing, does the organisation investigate or attribute blame?
- Can the business articulate what worked, what did not, and why — across multiple initiatives?
- Is there an explicit budget for experiments — work whose primary purpose is to generate learning?

Failure modes

Launch without defined success criteria. Measurement deferred or skipped. Results interpreted selectively to confirm prior commitments. Blame culture preventing honest assessment of failures. Learning retained within individual teams rather than shared as organisational capital.

Observable signal of health

Major initiatives carry explicit hypotheses. Results inform decisions about scaling, iterating, or stopping. Failures are analysed for systemic insight rather than individual accountability.

Table 4: Business Agility Diagnostic Summary

Dimension	Fitness Alignment	Key Diagnostic Question	Failure Mode
Portfolio Management	Fit for Purpose	Is work aligned to strategy and evaluated against ROI before commitment?	Fragmented resources, competing priorities, strategic incoherence
Effective Value Streams	Fit for Execution	Can strategy flow end-to-end with minimal coordination loss?	Handoff delays, invisible bottlenecks, slow strategic translation
Clear ROI	Fit for Improvement	Does execution generate learning that feeds back into strategy?	Repeated failures, unmeasured outcomes, strategy disconnected from reality

7. System Agility

System Agility is the emergent condition that determines whether the other three levels compound into durable organisational capability or remain locally functional but systemically brittle. It is not a level to be achieved after Business Agility; it is a continuous condition that either exists across the model or does not.

This distinction matters. A maturity model perspective would position System Agility as the highest stage — the most sophisticated form of agility. The 4LM positions it differently: as a meta-condition whose presence or absence shapes the quality of what is happening at every other level simultaneously. An organisation can have high-functioning teams, coherent structure, and responsive strategy and still lack System Agility if those levels are not reinforcing each other, if the organisation cannot assess its own fitness honestly, and if learning does not travel across boundaries.

7.1 Meta-Fitness — AFS Fitness Applied Reflexively

Meta-fitness is the capacity of the system to apply the AFS Fitness Sequence to itself. It is not enough for individual interventions to be assessed for fitness; the organisation's overall approach to its own development must also be fit for purpose, fit for context, fit for execution, and fit for improvement. Without this reflexive capacity, organisations optimise for the wrong things with sophistication, confuse performance for fitness, and mistake activity for progress.

Diagnostic questions

- Can the organisation distinguish between performing well and being fit? Does it interrogate the difference?
- When a change effort is launched, is the fitness of that effort itself assessed — not just its content?
- Is there a mechanism for questioning the organisation's current direction, not just its execution?
- Does the organisation recognise when it is optimising for a context that no longer exists?
- Are the criteria by which fitness is assessed themselves subject to revision?

Failure modes

Fitness conflated with performance metrics — an organisation that is performing well on its own measures while becoming less fit for its actual environment. No mechanism for questioning the validity of the current direction. Transformation efforts assessed by adoption of practices rather than by fitness outcomes. The organisation becomes increasingly sophisticated at doing the wrong things.

Observable signal of health

The organisation periodically steps back from execution to ask whether it is working on the right things, not just working on things correctly. Change efforts are assessed for fitness before being resourced. There is a recognised and used mechanism for strategic challenge that operates independently of hierarchy.

7.2 Cross-Level Coherence — The Four Levels as a Reinforcing System

Cross-level coherence is the condition in which the four levels of the model function as a reinforcing system rather than as separate improvement programmes. It is the organisational expression of strategic identity: the whole holding together under changing conditions, not because the parts are individually well-managed, but because they are aligned to a shared logic and mutually enabling.

Incoherence at this level is a specific failure mode distinct from failure at any individual level. An organisation can have high team capability, reasonable governance, and a defined portfolio process and still lack cross-level coherence if those elements are not reinforcing each other. Team autonomy without portfolio clarity generates well-executed work in the wrong direction. Portfolio clarity without execution capability generates strategic ambition that cannot be operationalised.

Diagnostic questions

- Do improvements at one level propagate visibly to other levels, or remain locally contained?
- Is there a shared logic — not just a shared vocabulary — that connects how decisions are made at the team, organisational, and business levels?
- When levels are in tension, is there a mechanism for resolving that tension coherently rather than through power or default?
- Does the organisation's strategic identity remain recognisable across conditions that change — restructuring, leadership transitions, market shifts?

- Are the feedback loops between levels operating — does what teams learn reach strategy; does strategic clarity reach teams?

Failure modes

Levels improving independently without compounding effect. Team-level agility gains neutralised by organisational governance that has not adapted. Business strategy refined without any connection to the execution capacity actually available. Transformation programmes running in parallel with contradictory logics. Strategic identity expressed in values documents but not in the operating model.

Observable signal of health

An improvement at the team level is visibly enabled by a prior change at the organisational level, which was itself enabled by a prior strategic clarification. The organisation can describe how its parts cohere — not as an aspiration but as an observable reality.

7.3 Adaptive Learning — Sense–Learn–Adapt at Enterprise Scale

Adaptive Learning is the enterprise-scale expression of the sense–learn–adapt loop that AFS's Adaptive Intelligence prescribes at every level. It is not team-level retrospection or business-level strategy review; it is the capacity of the system as a whole to process what it is learning — from markets, from execution, from failure, from experimentation — and adjust coherently across levels simultaneously.

The failure mode is not an absence of learning at any individual level. It is the absence of learning that travels. Most organisations generate significant local learning that never reaches the parts of the organisation that need it. Teams learn from incidents that never inform architectural decisions. Business units learn from market signals that never reach team prioritisation. The result is an organisation that is locally intelligent and systemically slow.

Diagnostic questions

- Is there a mechanism by which what teams learn from execution reaches strategic decision-making — and vice versa?
- When the organisation encounters a significant failure, does the learning from it become structural, or remain anecdotal?
- Are experiments designed to generate learning that will travel — or is experimentation conducted with no expectation that findings will inform other parts of the system?
- Is there a shared model of what the organisation is learning about itself — not just individual teams reporting metrics?
- Does the organisation have a recognised practice of questioning its own assumptions based on what execution reveals?

Failure modes

Learning isolated within teams or business units. Significant failures generating blame rather than structural insight. Experimentation conducted without sharing findings. No mechanism for market intelligence to reach team-level prioritisation. The organisation repeating failures at scale because the learning from their smaller-scale occurrence never propagated.

Observable signal of health

Learning from execution reaches strategy. Learning from strategy reaches execution. The organisation can articulate what it has learned about itself in the last quarter and how that has changed how it operates.

Table 5: System Agility Diagnostic Summary

Dimension	Fitness Alignment	Key Diagnostic Question	Failure Mode
Meta-Fitness	AFS Fitness (reflexive)	Can the system assess and question its own fitness, not just its performance?	Optimising the wrong things with sophistication; activity mistaken for progress
Cross-Level Coherence	FVI Coherence	Do the four levels reinforce each other and hold a recognisable strategic identity?	Local improvements that do not compound; levels pulling in different directions
Adaptive Learning	Adaptive Intelligence	Does learning travel across levels and inform the whole system?	Repeated failures at scale; isolated intelligence; systemic slowness

8. Using the Model: A Five-Phase Approach

The 4LM is a diagnostic and developmental tool, not an implementation roadmap. The five phases described here are not sequential steps to be completed once; they are a recurring practice to be applied as conditions evolve.

Phase 1 — Frame Through Fitness

Before assessing, clarify the engagement through AFS Fitness dimensions. This prevents the common failure of initiating a transformation effort whose purpose is assumed rather than tested.

- Purpose: What business problem is this organisation actually trying to solve through agility? What is the gap between current and desired capability?
- Context: What are the structural, cultural, and historical conditions in which this will operate? What will not change regardless of the transformation?
- Execution: What capabilities and resources actually exist? What is the realistic tolerance for disruption?
- Improvement: How will the organisation know this is working? What outcomes matter, and how will they be measured?

Phase 2 — Assess Current State

Use the four-level model to diagnose where coherence breaks down. For each level and each pillar, assess the current state against the diagnostic questions in Sections 4–7, identify the actual blockers, and map the dependencies between levels.

The assessment should be multi-method: metrics drawn from tooling and systems, observation of ceremonies and working practices, structured conversation with teams and leadership, and direct

experience of how work actually flows. Each method surfaces different kinds of evidence; none alone is sufficient.

Phase 3 — Apply Reverse Validation

Work backward from the desired outcome to surface hidden dependencies and identify the actual constraint. Start with Business Agility: if responsive, learning-oriented business agility is the goal, what must be true at the organisational level? Then work to Team Agility: if that organisational condition is required, what must be true at the team level?

This process identifies the critical path. The gap that blocks all others is the one to address first, not the most visible gap or the gap in the area where improvement is most politically palatable.

Phase 4 — Design and Sequence Interventions

Address the critical constraint first. Design interventions at that level and sequence the follow-on work that becomes possible once the constraint is removed. Build in quick wins alongside structural changes — not as distractions from the real work, but as the credibility-building that makes structural change politically viable.

The sequencing is not prescriptive. If Team Agility is the critical constraint, fixing it is the starting point. If Organisational Agility is the constraint, structural and governance changes take precedence regardless of how much more tractable team-level practice improvement might appear.

Phase 5 — Anchor Through Adaptive Intelligence

Establish the continuous sense–learn–adapt loop that prevents the model from being applied once and abandoned. This is not a monitoring programme; it is a practice of honest inquiry. What is actually improving? Has the context shifted in ways that require the approach to change? Are unintended consequences emerging? Is learning from execution reaching strategy?

The cadence is not fixed. Monthly reviews of metrics and progress. Quarterly retrospectives on the transformation itself. Regular structured listening to teams, leadership, and where possible customers. Incident and experiment reviews conducted with the explicit purpose of generating systemic insight rather than assigning accountability.

9. Coherence Matrix: Linking Fitness to Constraints

The Coherence Matrix below links each pillar across all four levels to its AFS Fitness dimension, the most common structural constraint that prevents it from functioning, and the Reverse Validation chain that shows how a failure at that pillar propagates through the model. It is intended as a diagnostic reference, not a checklist.

Table 6: Full Coherence Matrix — 4LM

Level	Pillar	Fitness Dimension	Common Constraint	Reverse Validation Link
Team Agility	Clarity of Intent	Fit for Purpose	Shifting priorities; no shared definition of done	Unclear intent → rework → broken cadence → no learning

Team Agility	Execution Capability	Fit for Execution	Monolithic architecture; no CI/CD; knowledge silos	Coupling → forced coordination → no independent action
Team Agility	Delivery Commitment	Fit for Improvement	Blame culture; retrospective theater	No safety → hidden problems → stagnation
Organisational Agility	Coherent Structure	Fit for Context	Functional silos; unclear ownership	Silos → dependency overhead → team autonomy impossible
Organisational Agility	Suitable Governance	Fit for Purpose	Approval bottlenecks; risk-theater governance	Slow decisions → wait time → execution decoupled from strategy
Organisational Agility	Meaningful Metrics	Fit for Improvement	Vanity metrics; no outcome visibility	Activity metrics → no insight → strategy drifts
Business Agility	Portfolio Management	Fit for Purpose	Strategy unclear; prioritisation by noise	No WHY → fragmented resources → incoherent execution
Business Agility	Effective Value Streams	Fit for Execution	Handoff delays; shared bottleneck services	Coordination required → slow delivery → market lag
Business Agility	Clear ROI	Fit for Improvement	Launch without success criteria; deferred measurement	No feedback → repeated mistakes → strategy disconnects from reality
System Agility	Meta-Fitness	AFS Fitness (meta)	No capacity for self-assessment; fitness conflated with performance	Can't assess own fitness → optimises wrong things → coherence erodes
System Agility	Cross-Level Coherence	FVI Coherence	Levels function in isolation; gains at one level don't propagate	Incoherence → local improvements don't compound → system stays brittle
System Agility	Adaptive Learning	Adaptive Intelligence	Learning isolated; no enterprise-scale sense–learn–adapt loop	No shared learning → same failures recur → adaptation is reactive not systemic

10. Illustrative Cases

The following cases are drawn from software development contexts and are presented to illustrate how the model's diagnostic logic operates in practice. They are composites; identifying details have been abstracted.

The Architecture Trap

A mid-market software organisation had developed mature team practices. Sprint cadences were reliable. Retrospectives generated and implemented changes. Teams could articulate what they were building and why.

Business agility, however, remained elusive. New features required months from decision to delivery. Strategic pivots were painful rather than routine.

The diagnostic revealed the constraint: a monolithic codebase that all teams touched. Every deployment required organisation-wide coordination. Team autonomy was effectively neutralised by an architectural decision made years earlier in a different strategic context.

The trap was the obvious intervention. The organisation had been investing in team-level process improvement — better backlog management, clearer acceptance criteria, more effective retrospectives. These were not wrong; they were irrelevant to the actual constraint. Organisational

structure — specifically, technical architecture — was preventing Business Agility regardless of how well teams executed.

Applying Reverse Validation revealed this clearly: if Business Agility requires value streams that can flow independently, what must be true about team autonomy? If team autonomy requires independent deployment, what must be true about architecture? The critical path ran through architectural decomposition, not sprint improvement.

After investment in service decomposition, deployment frequency increased tenfold within six months. Not because teams improved their practices — those were already mature — but because the organisational constraint was removed.

The lesson is structural: coherence requires addressing the actual constraint, which is frequently located at a different level than where the symptoms are most visible.

The Metrics Mirage

An enterprise organisation had coherent team practices and a well-governed organisational structure. The portfolio process was defined and followed. Value streams were mapped. Deployment frequency was high.

Business agility remained absent. Strategic priorities did not change in response to market signals. Resource allocation was stable in ways that had no relationship to what was working.

The diagnostic identified the failure mode: the organisation measured activity at every level and outcomes at none. Velocity, deployment frequency, and code coverage were tracked rigorously. User adoption, customer satisfaction, and revenue impact were not.

Without outcome metrics, the feedback loop from execution to strategy was broken. The organisation was generating significant learning at the team level — features that users did not adopt, experiments whose results were ambiguous — and none of it was reaching strategic decision-making. Portfolio priorities were set based on assumptions made at the beginning of the planning cycle and never revised.

This is a System Agility failure as much as a Business Agility failure. The Adaptive Learning pillar was absent: learning was being generated locally and not travelling. The result was an organisation that was operationally sophisticated and strategically blind.

The intervention required changes at two levels simultaneously: instrumentation for outcome measurement at the Business Agility level, and a structured mechanism for findings to reach portfolio decision-making. The two changes had to happen together; instrumentation without a receiving mechanism produces data that no one acts on.

11. Maturity Rubric

The rubric below provides a reference for assessing the current state of each pillar. It is a diagnostic instrument, not a target. Not all organisations need Level 5 maturity on all pillars; fitness requires matching maturity to the actual strategic context, not maximising it.

Level	Description
1 — Emerging	Practice is inconsistent; success depends on individual effort. No systematic measurement or feedback. Improvements are sporadic; regressions are common.

2 — Repeatable	Practice is consistent within teams; basic standardisation is emerging. Basic measurement exists and is used sporadically. Improvements happen and some persist.
3 — Defined	Practice is standardised across the organisation with conscious decisions about deviations. Metrics inform decisions regularly. Root causes of failures are addressed, not just symptoms.
4 — Managed	Practice is rigorous; measurement is continuous. Metrics drive explicit trade-off decisions. Improvements are data-driven; learning is organisational rather than local.
5 — Adaptive	Practice evolves continuously based on what the system is learning. Metrics reveal emerging patterns before they become constraints. The organisation adapts in real time; its learning is a strategic asset.

> A note on maturity targets
>
> *The rubric is a diagnostic tool, not a benchmark. An organisation operating in a stable, regulated context with low strategic uncertainty may be optimally fit at Level 3 on several pillars. Pursuing Level 5 would consume resources producing sophistication that the context does not require. Fitness, not maximisation, is the criterion.*

12. Discussion

The Incoherence Problem

The central argument of this paper is that most organisational agility failures are failures of coherence rather than failures of practice. Individual practices — team ceremonies, governance processes, portfolio management — are frequently well-designed and poorly situated. They are introduced into a system of other practices that do not reinforce them, at a level that is not the actual constraint, in a sequence that does not respect the dependencies between levels.

The 4LM does not prescribe practices. It provides the diagnostic structure to understand which practices are relevant at which level, in what sequence, and under what conditions. This is the contribution of grounding the model in AFS: the fitness question is always prior to the practice question. What are we actually trying to become fit for? Only once that is answered does the choice of practice become well-formed.

System Agility as Emergent Property

The decision to treat System Agility as a continuous emergent condition rather than as a fourth sequential level has theoretical and practical implications. Theoretically, it reflects a systemic rather than additive understanding of organisational capability: the whole is not the sum of its levels, and there is a quality of organisational functioning — coherent, self-assessing, learning — that only manifests when the levels are reinforcing each other rather than merely coexisting.

Practically, it means that System Agility cannot be addressed as a project. It cannot be installed. It is cultivated through the quality of the practices at the other three levels and through the explicit design of the feedback mechanisms that connect them. An organisation that wants System Agility must ask, at every intervention: does this strengthen or weaken the coherence between levels? Does this create or consume the conditions for learning to travel?

Reverse Validation as Diagnostic Discipline

The FVI Lens's Reverse Validation discipline has particular value in the 4LM context because the model's most common misuse is addressing the most visible symptom rather than the actual constraint. Teams are the most visible unit of organisational agility — their practices are observable, their outputs are measurable, their improvement is tractable. Organisational governance and business portfolio management are less visible and more politically complicated. The result is a systematic bias toward team-level intervention regardless of where the actual constraint is located.

Reverse Validation corrects for this bias by making the dependency chain explicit before interventions are designed. If the desired outcome is Business Agility, what must be true at the Organisational level for that to be achievable? If that organisational condition requires Team Agility of a particular kind, what must be true at the team level? The diagnostic starts from the outcome and works backward, rather than starting from the most tractable intervention and working forward.

Limitations

The 4LM is a diagnostic framework applied to a domain — organisational agility — that is inherently contextual. The model's pillars and diagnostic questions are drawn primarily from software development organisations. Applying the model to other domains requires adaptation of the diagnostic language, though the underlying logic — AFS Fitness alignment, sequential dependencies between levels, System Agility as emergent condition — is intended to be domain-agnostic.

The maturity rubric is descriptive rather than normative. It describes stages of capability development but does not prescribe which stage is appropriate for which organisation. Determining the appropriate maturity target requires the fitness assessment that the model describes but cannot itself perform.

The model does not address the political conditions under which transformation occurs. Structural constraints are diagnosable; political constraints — the distribution of power, the tolerance for challenge, the personal stakes of key actors — are not. The model assumes that the organisation in question has sufficient willingness to act on what the diagnostic reveals. Where that willingness does not exist, the diagnostic produces clarity that cannot be acted on, which is a different kind of problem requiring a different kind of intervention.

13. Conclusion

Sustainable organisational agility does not emerge from the adoption of practices. It emerges from coherent fitness across all four levels of the model — and from the continuous, active presence of the meta-capacity to assess that fitness, maintain coherence under pressure, and ensure that learning travels.

The Four-Level Model of Organisational Agility provides the diagnostic structure to understand where coherence is breaking down, which level holds the actual constraint, and what must be true at each level for the levels above it to function. Grounded in the Adaptive Fitness System and the FVI Lens, it does not ask organisations to adopt it; it asks them to use it to see more clearly what is actually preventing the agility they are pursuing.

The question the model ultimately poses is not "how agile are we?" It is: "fit for what, by what logic, and how do we know?" The answer to that question — honestly pursued and acted upon — is the beginning of System Agility.

References and Related Works

The following works constitute the foundational and predecessor documentation for this paper. All are published by Mario Aiello under the Adaptive Ways brand at adaptiveways.site.

Predecessor documents (superseded by this paper)

- Aiello, M. (2025). Software Development Agility Approach: Three-Level Model Grounded in AFS-FVI.
- Aiello, M. (2025). How the Three-Level Model Sits Within the AFS-FVI Ecosystem.
- Aiello, M. (2025). Software Development Agility Approach: Three-Level Model — A Practitioner's Guide.

AFS Foundational Works

- Aiello, M. (2024–2025). Adaptive Fitness System: White Paper Series. adaptiveways.site.
- Aiello, M. (2024–2025). AFS Logic: Four Sequencing Principles. adaptiveways.site.
- Aiello, M. (2024–2025). AFS Dynamics. adaptiveways.site.
- Aiello, M. (2024–2025). Fitness-Driven Design. adaptiveways.site.
- Aiello, M. (2024–2025). AFS Portfolio Management. adaptiveways.site.