WHITE PAPER

The Operating Model is the Bottleneck: Seven Principles for Product Delivery when AI is in the Loop

As AI makes building software faster and cheaper, this paper offers seven principles for product leaders, heads of engineering and delivery, CIOs and CTOs, Product Owners, Scrum Masters and coaches who want a clear view of how the operating model around their teams must change: what to keep, what to redesign and which signals to watch so that faster output becomes better outcomes.

Dave West

November 2026

©Value Amplified. All Rights Reserved.

CONTENTS

Table of contents

Table of contents 2
Summary 3
The question in front of every leader 3
Cheap building, expensive deciding 4
What the evidence shows 5
What this adds up to 7
Seven principles for the operating model 7
1. Start with the outcome 7
2. Own it end to end 8
3. Learn in small bets 9
4. Make decisions visible 9
5. Build governance into the work 10
6. Grow judgment, not just output 10
7. Evolve the model with evidence 11
What does not change 11
Two readings of a fast-moving situation 12
What would tell them apart? 13
What to do next 13
Appendix 15
Where did APOM's 59 principles go 15
Acknowledgments and disclosure 16
Sources 16
About the Author 17
About Value Amplified 17

©Value Amplified. All Rights Reserved

Summary

AI has made building software faster and less expensive. However, it has not made deciding what to build, checking whether it worked, or owning the result any cheaper. Most operating models were designed around the opposite assumption. Annual funding, approval gates, project charters, and handoffs all exist to ration scarce build capacity. When that constraint goes away, the machinery doesn't speed up. It either slows the gains down or waves more of the wrong work through.

This has five implications:

1. **The operating model, not the tooling, now decides what AI is worth.** AI amplifies whatever it lands in. A project-funded, gate-governed organization gets more output and the same outcomes.
2. **The operating model protects three things AI can't supply.** Direction: clear, measurable outcomes. Ownership: a named human accountable for every result. Evidence: short loops that show whether the work changed anything.
3. **Seven principles are enough.** They build on the Agile Product Operating Model, cut from 59 principles to a set that a leadership team can carry into a meeting.
4. **Signals, not scorecards.** Each principle comes with one or two signals to watch, not a full measurement system.
5. **Start where it hurts.** The principles are a diagnostic, not a rollout plan. Pick one, run a small experiment, check the signal.

This paper is for product leaders, heads of engineering and delivery, CIOs and CTOs, Product Owners, Scrum Masters and coaches who want a clear view of what the operating model around their teams must become. It's a companion to *From Production to Judgment: What AI Changes in Agile Product Delivery* (Alan W. Brown and Dave West, 2026), which looks inside the team. This paper looks at the system around it.

The question in front of every leader

Many leaders treat AI as a tooling decision. Pick the assistants, buy the licenses, train people on prompts, and the organization will speed up on its own. The instinct is understandable. That's roughly what happened with every previous wave of tooling and technology.

This time the tools are changing the economics that the operating model was built on. Google's DORA team put it plainly in its 2025 State of AI-assisted Software Development

©Value Amplified. All Rights Reserved

report: "AI's primary role in software development is that of an amplifier. It magnifies the strengths of high-performing organizations and the dysfunctions of struggling ones." The tools are already in use. DORA found that 90% of survey respondents use AI in their work. The question is what they're amplifying.

For most organizations, the answer is an operating model designed for a world where building was the slow part. Projects get funded once a year. Work passes from the people who decide, to the people who build, to the people who run it. Governance waits at the end of the line. Speed up the building inside that system, and you don't get a faster organization. You get a longer queue in front of the next gate.

This is a pro-product argument, not an anti-process one. Fund products, not projects; govern by outcomes; give long-lived teams real ownership; learn from evidence. AI doesn't make those ideas new. It makes them urgent.

Where this comes from. These principles build on the Agile Product Operating Model (APOM), which Dave West developed with colleagues and the community while CEO of Scrum.org. APOM belongs to Scrum.org, and this paper doesn't reproduce it. If you want the original, start with Project to Product and the Agile Product Operating Model and APOM Guiding Principles Have Published. APOM's guiding principles cover four areas, twelve sub-areas and 59 principles. That makes a strong reference, but it's hard to carry into a meeting. This paper cuts them to seven, asks what each must protect when humans and agents share the work, and gives each one or two signals to watch. The appendix shows where every original idea went.

Cheap building, expensive deciding

Every operating model encodes a belief about what's scarce. For most of the last fifty years, the scarce thing was building capacity. Developers were expensive, building took months, and a wrong bet cost a year. So organizations built machinery to ration it.

That approach reflected the constraints delivery professionals faced. A business case forced someone to say why the work mattered before anyone paid for it. A phase gate forced a pause to check before more money went in. A handoff to an operations team made sure someone would run what got built. The cost of building quietly enforced discipline. Nobody funded a team for a year without some kind of case.

AI removes that enforcement. In announcing AWS's *Reimagine* research, Tom Godden described what happens next: Building gets dramatically cheaper, "but the cost of deciding, funding, governing and approving it did not change." The constraint moves from building to everything around it.

©Value Amplified. All Rights Reserved

The mechanism	What is rationing	What it does when building is cheap
Annual project funding	Scarce build capacity	Locks in last year's bets while the cost of trying new one's collapses
Business cases and approval gates	The risk of an expensive wrong bet	Becomes the queue that absorbs the gains
Handoffs between deciding, building and running	Specialist time	Lets work reach production with no one feeling they made the call
End-of-line governance	Late, costly rework	Can't keep pace with work produced at machine speed
Output measures (velocity, utilization)	Productive use of expensive people	Rewards the one thing that's now cheap

So what is the operating model once output is significantly cheaper and faster? A redesigned operating model exists to protect three things AI can't supply on its own:

- **Direction.** Clear, measurable outcomes that tell people and agents what matters. When an agent will build whatever it's pointed at, the goal is your main control.
- **Ownership.** A named human is accountable for every outcome, regardless of who or what did the building. Accountability doesn't transfer to a tool.
- **Evidence.** Short loops that show whether the work changed anything. Fast output without fast learning is just faster inventory.

The seven principles discussed later each protect one of these. That's the test for anything else in your operating model too. If a practice doesn't protect direction, ownership, or evidence, ask why we still do it.

What the evidence shows

The evidence on AI in product delivery is moving fast, and most of it was gathered on tools that have since been overtaken. This paper relies on the direction it points rather than any single productivity or velocity number. The companion paper, *From Production to Judgment*, covers the task- and team-level evidence: Real gains on individual tasks, a growing review

©Value Amplified. All Rights Reserved

bottleneck, and the risk of capability debt. This paper focuses on what the evidence says about the system around the team. Three findings stand out.

Finding 1: AI amplifies the system it lands in

DORA's 2025 report found that AI adoption "now improves software delivery throughput, a key shift from last year. However, it still increases delivery instability." Its explanation is the important part: "while teams are adapting for speed, their underlying systems have not yet evolved to safely manage AI-accelerated development." More throughput plus more instability is what you'd expect when faster teams sit inside an unchanged operating model.

Finding 2: The gains are absorbed by everything around the work

DORA calls value stream management, "the practice of visualizing, analyzing, and improving the flow of work from idea to customer," a "force multiplier for AI." It's what ensures "that local productivity gains translate into measurable improvements in team and product performance." Without someone owning the whole flow, local gains leak away.

AWS's *Reimagine* research shows where they leak. In one AWS example, a rebuild scoped at thirty engineers for a year was done by five engineers in seventy days. Yet, in Rahul Pathak's words, "our security review, our deployment process, the UI review was still the bottleneck." The same research found that "funding processes are not built to reprioritize quickly." Approval, governance and funding were designed for a slower world, and they now set the pace.

DORA makes the product point directly about platforms: "Organizations that treat their platform as an internal product designed to improve developer experience see significantly greater returns." Product thinking isn't just for customer-facing work. It's how the operating model keeps up.

Finding 3: Most bets miss, so more bets need more evidence

Cheaper building doesn't make ideas better. In Online Experimentation at Microsoft (2009), Ronny Kohavi and colleagues reported that "only about 1/3 of ideas improve the metrics they were designed to improve." Those were ideas that had already passed Microsoft's internal reviews. If most bets miss, building more of them faster, with no measurable outcome to check against, just multiplies the misses. The organizations that benefit will be the ones that learn faster, not just build faster.

©Value Amplified. All Rights Reserved

What this adds up to

The pattern is consistent. AI speeds up the parts. Whether the whole improves depends on the operating model: How work is funded, who owns it end-to-end, where governance sits and how fast evidence turns into decisions. That's where the seven principles focus.

Seven principles for the operating model

Each principle protects one of the three things AI can't supply. They are deliberately few. A principle nobody can remember is a principle nobody uses.

#	Principle	In one line	Protects
1	Start with the outcome	No bet without a measurable why, for people or agents.	Direction
2	Own it end to end	One team owns discovery, delivery and support. A human owns every outcome.	Ownership
3	Learn in small bets	Every increment tests something. Output that teaches nothing is noise.	Evidence
4	Make decisions visible	Goals, assumptions and calls are visible, especially the ones AI makes easy to skip.	Direction
5	Build governance into the work	Guardrails are outcomes, built into the flow and inspected often.	Evidence
6	Grow judgment, not just output	Talent, incentives and AI use build capability instead of drawing it down.	Ownership
7	Evolve the model with evidence	The operating model is a product too: inspect and adapt it on a cadence.	Evidence

1. Start with the outcome

Every bet, whether it's a product strategy, a Product Goal, a Sprint Goal, or a prompt to an agent, starts with a measurable, why. If you can't say what will be different for a customer or the business, and how you'll know, you aren't ready to build. This was always good advice. Cost used to enforce it.

©Value Amplified. All Rights Reserved

What changes with AI. Cost no longer forces the question. An agent will happily build whatever it's pointed at. When output is nearly free, an unclear goal doesn't slow you down. It just produces more of the wrong thing, faster. Given that most ideas miss (Finding 3), the goal becomes the main control you have over what humans and agents spend their time on.

Signals to watch

- Share of work in progress, including agent tasks, that traces to a measurable outcome.
- Time from new evidence to a changed priority.

Try: Make the Product Goal the filter every backlog item and agent task must pass. Say no more often.

2. Own it end to end

One long-lived team owns a product across discovery, delivery and support. It talks directly to the people it serves and lives with what it ships. Handoffs between "the people who decide," "the people who build," "the people who use it," and "the people who run it" are where context, speed and accountability leak away.

What changes with AI. Agents add a new kind of handoff, and it's the easiest one to ignore. Work can go from a prompt to production with no one feeling they made the call. Every outcome, quality bar and decision to ship needs a named human owner. The team, not a tool, is still the unit of delivery. DORA's finding on value stream management (Finding 2) is end-to-end ownership by another name: someone has to see the whole flow, and own it.

Signals to watch

- Handoffs per change, from idea to running in production.
- Share of outcomes and agent workflows with a named, accountable owner.

Try: Map one product's flow from idea to customer. Count the handoffs and name the owner of each agent workflow.

©Value Amplified. All Rights Reserved

3. Learn in small bets

Work moves in small increments. Each one tests an assumption about value, usability, or feasibility, and is a chance to stop, pivot, or double down. A release you can't learn from isn't an increment. It's inventory.

What changes with AI. AI makes increments much cheaper to produce. It doesn't make them cheaper to validate. The bottleneck moves from building to learning: reviewing, testing with users and reading the results. Teams that ship more but learn at the same rate don't get faster. They pile up unchecked work. The companion paper presents evidence that how fast work *feels* is a poor guide to how fast it is.

Signals to watch

- Time from release to learning: How long until you know whether it worked.
- Share of increments whose result changed a decision about what to build, delay, or stop.

Try: For the next three releases, write down before shipping what you expect to learn and when you'll check.

4. Make decisions visible

Goals, assumptions, constraints, trade-offs and the reasons behind leadership calls are visible to everyone affected and open to challenge. Leaders set direction and guardrails, then let the people closest to the problem decide inside them. Transparency is what makes that delegation safe.

What changes with AI. A lot of organizational friction was quietly doing a job. Purchase orders, design reviews and estimation sessions forced people to stop and decide, together, out loud. AI removes much of that friction. The decisions still get made. They just get made silently, inside a prompt or a default. A decision nobody makes out loud is still a decision, just one nobody is accountable for. DORA lists a clear AI policy among the seven foundational capabilities that shape whether AI pays off. Knowing what's been decided is part of what makes AI work.

Signals to watch

- Decision latency: time from an issue being raised to a decision.
- Share of significant decisions recorded with their owner and key assumptions.

©Value Amplified. All Rights Reserved

Try: Keep a one-page decision log for a quarter: the call, the owner, the assumption it rests on, and when you'll revisit it.

5. Build governance into the work

Governance and compliance are stated as outcomes (what must be true), not procedures (who must sign). Teams choose how to meet them. Checks live in the flow of work and the Definition of Done, not in a gate at the end. Governance itself is inspected and adapted like anything else. The same goes for suppliers and partners: bring them in early, on contracts that can change as you learn.

What changes with AI. Phase-gate governance can't keep up with work produced at machine speed. Reviews pile up, or worse, get waved through. AWS's experience (Finding 2) shows the gate becoming the bottleneck even when the building is done. Short cycles, explicit goals and a real Definition of Done are the most practical guardrails for AI-assisted and agentic work. Empiricism is the governance model. It gives risk and compliance teams more evidence, more often.

Signals to watch

- Share of governance and compliance checks embedded or automated in the Definition of Done.
- How often governance effectiveness is reviewed, and what changed as a result.

Try: Pick one approval step. Rewrite it as the outcome it protects and move the check into the Definition of Done.

6. Grow judgment, not just output

The organization invests in its people's capability: Skills, mentoring and communities of practice that sit alongside product teams. Incentives reward outcomes, collaboration and learning, not volume or local optimization. Metrics are for improvement, never for ranking people.

What changes with AI. AI can raise output today while eroding the skills and judgment you'll need tomorrow. The companion paper calls this capability debt and sets out the evidence. The risk comes mostly from passive delegation, not from using AI itself. People who use AI to understand the work keep their skills; those who just hand the work off lose them. Incentives that reward output accelerate the debt. When building is cheap, judgment is the scarce asset. Treat it that way.

©Value Amplified. All Rights Reserved

Signals to watch

- Team confidence in explaining, changing and supporting the work AI helped produce.
- Whether rewards and promotion track outcomes and mentoring, or output volume.

Try: Remove output measures such as velocity from performance conversations. Decide on purpose how junior people will learn.

7. Evolve the model with evidence

The operating model is a product too. It has users (the teams), it has outcomes, and it should be inspected and adapted on a regular cadence. The same goes for how money flows: fund long-lived products and steer them often, rather than funding projects once and hoping. Change isn't a transformation program with an end date. It's how the model works.

What changes with AI. The tools change every few months. An operating model designed once, for this year's AI, will be out of date before the rollout finishes. Funding that can't reprioritize quickly (Finding 2) turns cheap experiments into slow ones. Teams need to own and adapt how they work, including how they use agents, and leaders need a cadence for changing the wider model as evidence comes in.

Signals to watch

- Time from an improvement being identified to it being in place.
- How often investment and operating-model decisions are revisited based on delivered outcomes.

Try: Move one product from annual to quarterly funding review, steered by its outcomes rather than its plan.

What does not change

It's easy to read the last section as a list of things product operating models got wrong. It isn't. Every principle points the same way: Away from rationing, build capacity and toward protecting direction, ownership and evidence. That's where the product movement started. AI hasn't changed the destination. It has removed the excuse for not getting there.

Three things hold.

- **Empiricism.** Transparency, inspection and adaptation were never about software. They're how you make decisions under uncertainty. AI adds uncertainty, so empiricism matters more, not less.

©Value Amplified. All Rights Reserved

- **The product as the unit of investment.** A product has customers, outcomes and a lifecycle. A project has a budget and an end date. When the cost of building collapses, the thing worth funding is the outcome, and products are how you hold outcomes over time.
- **The team as the unit of delivery.** Individuals and agents can produce more than ever. Turning that output into value still takes a group of people who share a goal, challenge each other and live with what they ship. The team is the integration layer.

The hardest parts of product delivery have always been human: choosing which problem matters, having the courage to stop work the evidence says is failing, and getting an organization to absorb change. For decades they were partly hidden behind the cost of building. As building gets cheap, they're in plain sight. AI can't make those calls for you. A good operating model makes sure someone does.

Two readings of a fast-moving situation

The evidence supports the diagnosis. It doesn't yet say how far the operating model has to change. Experienced leaders looking at the same data reach different conclusions, and with the tools changing every few months, both readings are credible. A leadership team's choice between them is a real decision, so we lay both out side by side.

Reading one: Evolve the model you have

On this reading, the core of a product operating model is already right. Long-lived product teams, outcome-based funding, governance built into the flow and leaders who set direction rather than approve tasks are what AI-era delivery needs. The job is to finish the move from project to product that many organizations started and stalled, and to extend each element for agents: Owners for agent workflows, AI checks in the Definition of Done, faster funding reviews. The seven principles work as a checklist for that evolution.

Reading two: The structure around the team gets redesigned

On this reading, the principles hold, but the structures that carry them may not. If small teams with AI can do what large programs did, then team size, portfolio shape, funding units and the economics of scaling could all change more than incremental reform allows. AWS's Tom Godden writes that "no one knows yet what will replace the two-pizza team." Some organizations are experimenting at the edges. Best Buy, he reports, "deliberately kept its AI teams outside its standard delivery and review process," then brought the work back in for product review and brand guidelines.

©Value Amplified. All Rights Reserved

What would tell them apart?

The evidence to settle this doesn't exist yet, but leaders can watch for it. If organizations that evolve their existing product model sustain quality, outcomes and learning as agentic tools spread, reading one gains ground. If organizations that redesign funding and structure around much smaller units consistently do better, reading two does. The Best Buy example hints at a third possibility: the old process gets bypassed first and rebuilt later, which is costly either way.

Whichever reading you hold, the test for any part of the operating model is the same. Keep what protects direction, ownership and evidence. Be willing to let go of what only existed to ration building.

What to do next

Approach the update of the operating model as a complex problem that requires an incremental, evidence-based approach to improvement. Treat the principles as a diagnostic: Find the one where it hurts most, run a small experiment and check the signal.

Start with a one-hour diagnostic

1. **Score each principle from 1 to 5** with a product team and its key stakeholders. Do it together, in the open (see principle 4).
2. **Pick the lowest score that matters most** to the outcome you're chasing this quarter.
3. **Choose one change and one signal**, and agree when you'll look at the result.
4. **Inspect and adapt** on that date. Then pick the next one.

Executives and portfolio leaders

1. Stop treating AI tool adoption as the goal. Ask where the time saved is going, and what's absorbing it.
2. Move at least one product from annual project funding to outcome-based funding with a short review cycle.
3. Make one governance requirement an outcome, not a sign-off, and ask risk and compliance to help design the check.

©Value Amplified. All Rights Reserved

Heads of product and engineering

1. Map the flow from idea to customer for one product and count the handoffs, human and agent.
2. Name an accountable owner for every agent workflow that touches production.
3. Take output measures out of performance conversations. Measure outcomes, quality and capability.

Product Owners and product managers

1. Use the Product Goal to say no more often. A list of work can now be generated endlessly; a clear outcome can't.
2. Write down what each release should teach you, and check.
3. Track adoption of what ships, not just delivery of it.

Scrum Masters and agile coaches

1. Help leaders make decisions visible: who decided, on what assumption, and when it will be revisited.
2. Coach the organization, not just the team. The constraint has moved outside the Sprint.
3. Make the team's use of AI, and what it's doing to their skills, a standing Retrospective topic.

Most organizations will find the same thing. AI has already sped up the building. The rest of the operating model hasn't caught up.

AI changes how fast organizations can build. The operating model decides whether they still know why.

©Value Amplified. All Rights Reserved

Appendix

Where did APOM's 59 principles go

Every APOM sub-area maps to at least one of the seven principles, so nothing was dropped, only merged. Principles 3 and 5 absorbed the most, each drawing on six sub-areas.

APOM area	Sub-area	APOM principles	Now covered by
Strategy	Strategy	8	1, 2, 5, 7
People	Organization	5	2, 4, 6
People	Leadership	5	1, 4, 6, 7
People	Culture	6	1, 3, 4, 7
People	Talent and Skills	3	6
People	Incentives	5	4, 6
Structure	Governance and Compliance	3	5
Structure	Process	4	3, 5, 6, 7
Structure	Procurement	5	3, 5
Value Cycle	Discovery	5	1, 2, 3
Value Cycle	Delivery	5	1, 3, 5
Value Cycle	Support	5	2, 3, 5

The APOM approach paper's Agile Product Portfolio Management and Change Management sections map to principle 7.

©Value Amplified. All Rights Reserved

Acknowledgments and disclosure

APOM was developed at Scrum.org with Krystian Kaczor, Andy Brandt, Bogdan Onyschenko, Konstantin Razumovski, Chad Beier, Ziryan Salayi, Yuval Yeret, Olivier Ledru, Professional Scrum Trainers and the wider Scrum.org community. This paper builds on their work and doesn't reproduce it. The author, Dave West, was CEO of Scrum.org from 2015 to 2026 and now runs Value Amplified, which advises organizations on product delivery. Claude, made by Anthropic, helped research and draft this paper.

Sources

We checked quotations against the original pages on November 3, 2026.

- DORA (Google Cloud), 2025 State of AI-assisted Software Development: the amplifier finding (p. 3); 90% of respondents using AI, throughput and instability, value stream management as a force multiplier, platforms as internal products, seven foundational capabilities including a clear AI policy (pp. 4–5).
- Godden, T., Announcing the new AWS Reimagine report on AI, AWS Enterprise Strategy blog (September 24, 2026): the cost of deciding, funding, governing and approving; the AWS rebuild and Rahul Pathak quote; funding processes; the two-pizza team; Best Buy.
- Kohavi, Crook, Longbotham et al., Online Experimentation at Microsoft (2009): about one-third of ideas improve their target metric (p. 8).
- Brown, A. W. and West, D., *From Production to Judgment: What AI Changes in Agile Product Delivery* (2026): companion paper on task- and team-level evidence and capability debt.
- Scrum.org, The Agile Product Operating Model, Project to Product and the Agile Product Operating Model and APOM Guiding Principles Have Published: the original work this paper builds on.

©Value Amplified. All Rights Reserved

About the Author



Dave West is Partner and Founder of Value Amplified. His current research focuses on how agentic systems are changing product delivery, from team design and operating models to how value is measured. For more than a decade, Dave was CEO and Product Owner of Scrum.org, and before that Chief Product Officer at Tasktop, where he helped grow the company from a services business into a VC-backed product company. He previously led the software development and delivery practice at Forrester Research as VP and Research Director, ran Ivar Jacobson Consulting in North America, and led development of the Rational Unified Process (RUP) at IBM/Rational. A frequent keynote speaker, he is co-author of The Nexus Framework for Scaling Scrum and Head First Object-Oriented Analysis and Design.
LinkedIn

About Value Amplified

Value Amplified is a research-based advisory organization designed to help organizations, teams, and individuals to amplify their value. We empower teams to revolutionize how they design, develop, manage, and market digital products within an AI-driven market.

Contact us for more information.

©Value Amplified. All Rights Reserved