

Rewiring the Role

How AI Is Changing Coordination Work in Software Teams

Why product managers, scrum masters, and project managers who use AI to do the same job faster are solving the wrong problem.

The horseless carriage problem.

Most knowledge workers are using AI to do the same job, faster. A smaller number are using AI to do a different job, better. The second group is about to pull dramatically ahead of the first — and many in the first group don't yet realize the gap is forming.

This is the second paper in Kollabio's *Ground Truth* series on AI's impact on software organizations. The first — *Rewriting the Build* — covered changes across the software development lifecycle. This paper covers the humans who orchestrate that lifecycle: product managers, scrum masters, project managers, release train engineers, business analysts, and the wider coordination layer of modern product and engineering teams.

The pattern we observe across client engagements is consistent. Practitioners adopt AI tools, use them for a narrow slice of novel work, and then return to their traditional workflow for the bulk of their week. The Scrum Master who uses AI to generate epics from a feature description, but still writes status reports by hand. The PM who uses AI to summarize user interviews, but still manually constructs the quarterly roadmap deck. The project manager who uses AI to draft meeting notes, but still maintains the RAID log by hand. **This is the horseless carriage** — the new tool, wrapped around the old workflow.

Inside this paper:

- **The two-layer model of coordination work** — why most of these roles have always contained a mechanical "movement" layer and a human "judgment" layer, and why AI collapses only the first.
- **Ceremony-by-ceremony walkthrough** — a detailed treatment of backlog grooming, sprint planning, standups, DoR/DoD verification, status reports, risk management, sprint demos, retrospectives, and close-out. What changes, what stays human, what practitioners should do with the time they recover.
- **The rebalancing framework** — a practical model for how a coordination practitioner should restructure their week. From 70/30 movement-to-judgment toward 20/70/10 with a new slice for deliberate learning.
- **Federal and regulated environments** — how the ceremonies and roles translate into GovCon and regulated-industry contexts, including EVM, IMS, QASP, and FedRAMP-adjacent considerations.
- **The uncomfortable conclusion** — coordination roles are splitting into two paths. One is becoming obsolete. The other is becoming irreplaceable. Practitioners and leaders who don't recognize the split will be caught by it.

CONTENTS

What's inside.

01 **The Horseless Carriage**

A Scrum Master, an epic generator, and a weekly status report. The pattern in miniature.

02 **Coordination Work Has Always Been Two Layers**

Movement and judgment. Why AI collapses one and amplifies the other.

03 **Ceremony by Ceremony**

Backlog grooming, sprint planning, standups, DoR/DoD, status, risk, demo, retro, close-out.

04 **The Rebalancing Framework**

From 70/30 to 20/70/10. What a coordination practitioner's week should actually look like.

05 **Considerations for Federal & Regulated Environments**

EVM, IMS, QASP, compliance ceremonies, and the specifics of GovCon delivery.

06 **The Uncomfortable Conclusion**

Roles are splitting. One path is becoming obsolete. The other, irreplaceable.

07 **For Practitioners and Leaders**

Specific moves to make in the next 90 days, and the next 24 months.

ABOUT THE GROUND TRUTH SERIES.

*Ground Truth is Kollabio's series of field-level papers on AI's real impact inside software organizations. Paper 1 — **Rewriting the Build** — examined how AI is changing software development end-to-end, from Product Management through Deployment and Operations. Paper 2 — this document — examines how AI is changing the coordination work around that lifecycle. Papers 3 and 4 are in development. Each paper stands alone; together, they reflect Kollabio's view of how modern software organizations change in practice, not theory.*

A Scrum Master, an epic generator, and a weekly status report.

A Scrum Master we work with recently started using AI to generate epics and stories from a feature description. It's a good move — saves maybe thirty minutes per feature, produces better-structured backlog items than a tired human would write at 4 p.m. on a Thursday.

Then the same Scrum Master goes back to preparing the weekly status report by hand. Pulls numbers from Jira. Formats the table. Drafts the narrative. Pastes it into Confluence. Two hours, every week, forever.

This is the horseless carriage.

When cars first appeared, they were designed like carriages with the horses removed. It took a generation before anyone redesigned around what the new thing actually enabled. Most knowledge workers today — especially in coordination roles — are at that same stage. They're using AI to make existing artifacts faster instead of asking whether the artifacts should exist at all.

The epic generation saved thirty minutes. The status report wastes two hours. One was novel and felt permitted. The other was routine and felt like "the job."

That's the pattern this paper is about.

The diagnostic question for every practitioner reading this paper: *which of your weekly activities have you looked at with fresh eyes in the last six months?* The activities that feel most like "the job" — the status reports, the standups, the RAID logs — are exactly the ones most likely to be due for re-examination. The fact that they feel permanent is the signal, not the counter-argument.

We want to be specific about what's happening and what isn't. This is not a paper about AI replacing coordination roles. We do not believe that is happening, and we do not believe it is likely to happen in any near-term horizon we can reason about. The humans in coordination roles do things AI cannot do, and those things are not going away.

What is happening is that the composition of these roles is changing rapidly, and most practitioners are not rebalancing to match. They are using AI to accelerate the parts of the job that are becoming less valuable, while underinvesting in the parts that are becoming more valuable. Over a 24-month horizon, that creates a growing and visible gap between practitioners who rebalanced and practitioners who didn't.

The rest of this paper walks through what the rebalancing looks like, ceremony by ceremony, with a deep focus on the Scrum Master role — because the SM's ceremonies are the most codified and the contrast is sharpest there. The same pattern applies, with adjustments, to every coordination role in a modern software organization.

Coordination work has always been two layers.

Every coordination role — Product Manager, Scrum Master, project manager, business analyst, release train engineer — has always contained two distinct layers of work:

The movement layer

Information moving from one place to another. Jira tickets becoming sprint reports. Standup notes becoming leadership updates. Commits becoming burndown charts. Test results becoming release readiness memos. Backlog items becoming refined stories. This is mechanical work — important, but fundamentally about transit. The information exists; the role moves it from one shape or location to another.

The judgment layer

Deciding what matters. Deciding what to escalate and what to absorb. Reading the room during a steering committee meeting. Knowing which stakeholder needs a phone call before the meeting. Recognizing that the team's slowdown isn't capacity — it's morale. Figuring out how to say the thing that needs saying, in the room where it needs saying, in the way that lands. This is human work, and it is what coordination roles are actually for.

The pre-AI baseline

In a pre-AI world, the movement layer consumed 60–70% of a coordination practitioner's time in a typical week. Status reports, backlog refinement, RAID log maintenance, dashboard updates, meeting notes, and the dozens of small information-movement tasks that fill a calendar. The judgment layer — the actually valuable part of the role — got what was left after the artifacts were produced.

Most coordination practitioners we know have always found this ratio frustrating. The best ones carve out time to do judgment work anyway, often after hours or in the margins between meetings. The average ones produce adequate artifacts on

schedule and let the judgment work stay implicit. In either case, the shape of the role was defined by the movement layer.

What AI changes

AI collapses the cost of the movement layer toward zero.

Not to zero — some human oversight always remains, and should — but to a fraction of what it was. The mechanics of pulling data, formatting it, drafting narrative, distributing it, updating registers, generating summaries, and cross-referencing documents are exactly the profile of work AI handles well.

This is the good news. The bad news is what practitioners are doing with the recovered time:

Most practitioners are not using that recovered time to strengthen their judgment layer. They are using AI to do the movement work 30% faster, then filling the recovered 30% with more movement work.

More polished reports. More detailed dashboards. More elaborate RAID logs. More granular burndown analysis. Deeper Jira hygiene. The horseless carriage, extended to fill the available space. The ratio of movement to judgment stays roughly the same; the absolute volume of artifacts increases.

This is the pattern the rest of the paper is trying to disrupt. The rebalancing that should be happening — and isn't, for most practitioners — is a structural shift in how the week is spent, not an efficiency gain within the existing shape.

What changes, and what stays human.

This is the core of the paper. We walk through the major ceremonies of a modern software team in order, asking for each one: what was it designed for, what does AI change, and what should the practitioner be doing with the recovered time?

We focus on the Scrum Master role because its ceremonies are the most codified — the contrast between old and new is sharpest there. The Product Manager, project manager, and release train engineer roles follow the same pattern with discipline-specific adjustments. Where those adjustments are meaningful, we note them.

CEREMONY · 01

Backlog grooming

OLD WAY

The team reads through every story on the grooming list. Half the time is spent on mechanical readiness — does this story have acceptance criteria, are dependencies mapped, is it sized sensibly. The other half is spent on genuine judgment calls. Grooming sessions consistently run long, and half the stories get deferred to "next time."

AI-ENABLED

Every story is pre-checked against DoR criteria before the session. AI flags missing acceptance criteria, unmapped dependencies, vague language, and estimates that diverge from historical patterns for similar work. The team walks in with a pre-sorted list: the 80% mechanically ready and needing brief confirmation, and the 20% with actual judgment calls.

Stays human: the judgment calls themselves. The trade-off conversations. The "do we actually need this" question that usually got skipped for time. The team's understanding of how this story fits the product's direction.

Sprint planning

OLD WAY

The team assembles stories, estimates capacity, commits to a sprint goal, and argues about what fits. The mechanics of assembly — pulling velocity history, mapping dependencies, surfacing candidate stories — consume most of the session. The actual strategic conversation about what to commit to gets compressed into the final fifteen minutes.

AI-ENABLED

AI arrives at planning with a proposed sprint shape based on historical velocity patterns, story dependencies, and team capacity. The team's job is to validate and adjust rather than construct from scratch. Two hours of assembly becomes forty minutes of validation plus an hour of actual strategic conversation about the sprint goal.

Stays human: the commitment. The sprint goal. The conversation about what success looks like. The judgment about whether the proposed shape is right, not just whether it's feasible. The Scrum Master's reading of team morale and whether this is an ambitious sprint or a reasonable one.

Daily standup

OLD WAY

Each team member reports yesterday, today, and blockers. Fifteen minutes, every day, for a team of eight. Roughly 100 person-hours per year, per team, spent on what is primarily a status-reporting ritual. Much of what's shared could be read from the tools; the rest is the valuable part.

AI-ENABLED

Activity is already visible through AI-aggregated signals — commits, ticket movement, PR activity, CI status. The status portion of the standup is redundant. The remaining purpose is surfacing coordination needs ("I'm stuck and need help"), human-only context ("I'm worried about the approach"), and decisions ("we need a call on this today"). That's a five-minute conversation, not fifteen.

Stays human: the coordination asks. The worries. The decisions. The social fabric of a team that sees each other daily. These are worth preserving. The mechanical status reporting is not. Many mature teams run standups asynchronously via chat; others keep a short synchronous ritual. Either works. What doesn't work is running the old format as theater while AI makes the reporting half of it redundant.

CEREMONY · 04

Definition of Ready & Definition of Done verification

OLD WAY

A human checks every story against a DoR checklist at the start of development, and a DoD checklist at the end. Tedious, time-consuming, and — because it's tedious — often skipped, rubber-stamped, or applied inconsistently. The checklists become aspirational posters rather than enforced gates.

AI-ENABLED

AI runs the full checklist against every story automatically. Acceptance criteria present and testable? Dependencies resolved? Tests written? Docs updated? Accessibility verified? Security review complete? The practitioner reviews exceptions and edge cases, not every item.

Stays human: the judgment calls on edge cases. The decision about whether an exception is acceptable or a blocker. The cultural work of actually holding the line on quality gates when pressure to ship arrives.

CEREMONY · 05

Status reports

OLD WAY

The practitioner spends two hours a week pulling data, formatting tables, writing narratives, and distributing the result. Goes out Friday afternoon. By Monday, it's already stale. Readers skim or skip it. The exercise persists because it has always existed.

AI-ENABLED

Status is ambient and current. A leadership-ready view of program state is available on demand, generated from source systems and refreshed continuously. The practitioner stops *producing* the document and starts *curating* it — interpreting what the data means, flagging escalations, framing trade-offs for leaders who can't read the dashboard every day.

Stays human: the interpretation. The escalation judgment. The stakeholder management. The narrative that connects this week's data to the strategic arc of the program. A senior leader doesn't need a report; they need a practitioner who can tell them what matters and why.

Risk management

OLD WAY

A RAID log or risk register updated weekly during a manual review. Risks age, get stale, and often reflect what was worrying three quarters ago rather than what's worrying now. The log exists primarily to satisfy governance requirements.

AI-ENABLED

A living risk view maintained continuously from ongoing signals — slipping stories, failed builds, recurring defects, sentiment drift in retros, dependencies stuck in other teams' backlogs, aging PRs, outdated dependencies. The practitioner reviews, prioritizes, and drives mitigation, but doesn't maintain the register by hand.

Stays human: severity assessment. Mitigation strategy. The political work of getting other teams to unblock dependencies. The read on whether a risk is actually managed or whether the mitigation is theater. The register becomes a conversation-starter rather than a document-maintenance exercise.

Sprint demo / review

OLD WAY

The PM or SM assembles a demo narrative from what shipped. Slides, screenshots, talking points. Each stakeholder segment gets a framing. Hours of preparation for a thirty-minute meeting.

AI-ENABLED

AI drafts the narrative structure from what actually shipped — commit history, closed tickets, feature flags turned on. The practitioner's job is to tell the story of *why it matters* — connecting the work to user outcomes, business value, and strategic direction. The assembly is done for you; the meaning-making is yours.

Stays human: the story. Always the story. What did we learn this sprint? What hypothesis did we prove or disprove? What should stakeholders understand about where we are and where we're going? These are not in the commit log.

CEREMONY · 08

Retrospective

OLD WAY

The facilitator arrives with a blank template. The team has to remember what happened three sprints ago. Patterns across sprints go unnoticed because no one is looking for them. The retro often becomes a ritual complaint session without a clear commitment to change.

AI-ENABLED

The retro arrives pre-prepared with AI-surfaced patterns across multiple sprints — cycle time trends, defect clusters, recurring blockers, sentiment shifts, workflow anomalies. Humans bring opinions, stories, and emotional truth; AI provides the data backbone that makes those opinions actionable.

Stays human: the conversation. The courage to name what's actually wrong. The commitment to change something. The psychological safety work that makes honest retros possible. None of that is in the data.

CEREMONY · 09

Close-out and acceptance

OLD WAY

End-of-project documentation, lessons-learned reports, acceptance sign-offs, archival. The practitioner spends the final week of the project producing artifacts that almost no one reads. Closeout becomes a second shadow project that competes with actual delivery.

AI-ENABLED

Most documentation is generated from the actual history of the project — decisions, pivots, deliverables, test results. The practitioner reviews, edits, and adds the judgment layer: what actually matters for the next team, what the organization should learn, what customer-facing narrative closes the relationship well.

Stays human: the meaning-making. The customer conversation. The honest assessment of what went well and what didn't. The handoff to whoever comes next.

The pattern across all nine

If you read the ceremony cards above with fresh eyes, the same pattern repeats in every single one. The movement work compresses. The judgment work remains — and in most cases, becomes *harder* because it's no longer diluted by mechanical filler. The practitioner who embraces the shift does less assembly and more thinking. The practitioner who doesn't is increasingly doing work that AI can do, in front of colleagues who can see the difference.

From 70/30 to 20/70/10.

Walk through the nine ceremonies above and add up the time. A typical Scrum Master on a mature team spends something like 60–70% of their working hours on activities AI can substantially automate. A Product Manager, 50–60%. A project manager, 70–80%. If AI actually absorbed that mechanical load, practitioners would have 15–25 recovered hours per week — a full additional half-role's worth of capacity.

Almost no one we work with is using it that way. Instead, the recovered time is being absorbed by: more detailed versions of the same artifacts, additional meetings no one actually needs, dashboard proliferation, and the same work done more thoroughly in the same pattern. The tooling changed; the job didn't. The people doing the job mostly don't notice, because the work still feels like work.

Here is what the rebalancing should look like:

BEFORE — TODAY	AFTER — REBALANCED
70 / 30	20 / 70 / 10
Movement (artifacts, reports, ceremonies) 70%	Movement (AI-assisted, curated) 20%
Judgment (strategy, stakeholders, dynamics) 30%	Judgment (expanded, protected) 70%
Learning (deliberate skill building) 0%	Learning (deliberate, reserved) 10%

What each slice means in practice

Movement — 20%. The residual movement work that AI can't fully absorb, or where human review remains essential — exception handling in DoD checks, curation of AI-generated artifacts, edit passes on status narratives, review of AI-generated risk

assessments. The goal is not zero movement work; it is minimal, supervised, high-leverage movement work.

Judgment — 70%. The work that was always supposed to be the role. Strategic framing with leadership. One-on-ones with team members. Cross-team negotiation. Stakeholder management in its harder forms. The conversations that move organizations. The reading of team dynamics. The decisions about what to escalate, what to absorb, what to push back on, what to accept. For most coordination practitioners, this slice has been chronically under-invested for their entire career because the movement layer consumed the calendar. AI makes it possible to invest here; the rebalancing is what makes it real.

Learning — 10%. Deliberate, reserved time for learning. Not reading articles at lunch; actual time on the calendar to evaluate new tools, test new practices, and keep up with what's changing. In a world where tooling and practices shift quarterly, practitioners who don't protect this slice fall behind within twelve months — and the fall-off is visible to peers and leaders well before it's visible to the practitioner themselves.

A practical move for this week: audit your last two weeks of calendar and Jira activity. Count the hours spent producing artifacts vs. hours spent in strategic conversations or making difficult calls. If the first number is larger than the second, you are in the 70/30 pattern. The rebalancing is not a mindset shift; it is a calendar shift. It requires saying no to meetings and artifacts that used to feel mandatory, and saying yes to work that feels less visible but matters more.

Specific to GovCon and regulated delivery.

Coordination practitioners in federal and regulated environments face a set of specific constraints that shape how these rebalancing moves translate. The ceremonies above all apply; the overlays need attention.

Earned Value Management and Integrated Master Schedule

EVM reporting and IMS maintenance have historically consumed 15–25% of a GovCon project manager's time. The data-movement parts of EVM — CPI and SPI calculation, variance analysis, format compliance — are nearly textbook cases for AI automation. What remains is the judgment layer: explaining variances to program leadership, making the case for replans, framing recovery strategies. Practitioners who rebalance here recover significant capacity. Practitioners who continue to maintain EVM narratives by hand are increasingly doing work their commercial peers have already offloaded.

QASP and contract deliverable review

Quality Assurance Surveillance Plans and contract deliverable review cycles have ritualized the movement layer. Deliverable package assembly, compliance checks against DIDs, formatting against government templates — all of this is automatable. The human work that remains: the government liaison conversations, the negotiation when a deliverable doesn't land cleanly, the trust-building with a COR or COTR that makes difficult conversations possible when they need to happen.

ATO and compliance ceremonies

Authority to Operate processes and their adjacent compliance artifacts — POA&Ms, SSPs, continuous monitoring reports — are some of the most document-heavy coordination work in federal delivery. AI meaningfully accelerates the drafting, cross-referencing, and maintenance of these artifacts. It does not replace the judgment calls about risk acceptance, the stakeholder management with authorizing officials, or the navigation of agency-specific interpretations of NIST controls. Practitioners who

treat ATO as an automation opportunity *and* as a judgment discipline move faster than peers who treat it as only one of the two.

Labor category constraints

Federal contracting often constrains what labor categories can perform what work. As AI changes what work a given role actually does, practitioners and primes need to be deliberate about how role descriptions and labor category mappings are represented in proposals and delivery plans. This is not mature practice yet. Organizations getting it right are building language about AI-augmented delivery into their proposals and their internal role definitions, rather than pretending the work is unchanged.

Section 508 and accessibility ceremonies

Accessibility reviews are a canonical example of work AI accelerates but cannot complete. Automated scanning catches roughly 30–40% of Section 508 issues. The remaining 60–70% requires human judgment, often with assistive technology in hand. Practitioners who use AI to find the clear-cut cases quickly, and reserve their time for the judgment-required cases, do dramatically better accessibility work than peers doing it the old way.

The procurement-side implication

GSA MAS, BPAs, task order vehicles — the procurement apparatus of federal delivery — is slowly adjusting to AI-enabled delivery. Practitioners who can articulate, in proposal and briefing language, how their team uses AI to deliver better outcomes have a growing advantage over practitioners who cannot. This is a writing skill, a thinking skill, and an experience skill. It is also, increasingly, a competitive differentiator in GovCon capture.

Net effect for GovCon practitioners: the rebalancing framework from Section 4 applies with more force in regulated environments, not less. The movement layer is larger (because compliance is document-heavy), which means the recovery from AI-assisted rebalancing is larger. The judgment layer is also larger (because regulated delivery has more stakeholders and more consequence), which means the redirected capacity has more places to go. Practitioners who treat federal delivery as an exception to the AI-driven shift are misreading the environment. It is, if anything, an accelerator.

Coordination roles are splitting.

Coordination roles are not being eliminated. They are *splitting* — into two distinct career paths that look similar on the org chart and are diverging rapidly in what they actually are.

PATH A

Artifact-heavy

Produces documents. Runs ceremonies as rituals. Measures its value by the volume and polish of its outputs. Defines the role by the artifacts produced, not the decisions enabled.

Not unskilled — many on this path are highly skilled. But the work they do is work AI increasingly does better, faster, and at lower cost.

Becoming obsolete.

PATH B

Judgment-heavy

Anchors in the things AI cannot do: strategic framing, stakeholder navigation, team dynamics, difficult conversations, organizational politics, the creative reframing of problems.

The work the organization actually needs someone to do. The work that is harder to measure, harder to hire for, and more valuable than ever.

Becoming irreplaceable.

This split is already visible inside organizations that are paying attention. It is not yet visible inside organizations that are not — but it will be, within 24 months, when the next round of cost-cutting lands with noticeably different criteria than the last round. Coordination headcount will not be reduced uniformly. It will be reduced along the artifact/judgment axis, whether or not the organizations making those decisions articulate it that way.

What this means, directly

The question is not whether to adopt AI. You already have — or if you haven't, the people around you have. Organizational AI adoption is not a policy decision; it is a ground-level behavioral fact.

The question is whether you are using it to protect your existing role, or to transform it. Those are very different strategies, and they lead to very different places. *Protection* means: learning the tools well enough to keep doing the same job faster.

Transformation means: restructuring the week around the judgment layer and letting AI handle the rest.

Protection buys about 24 months. During that window, tool adoption continues to compound. Organizations that were tolerating artifact-heavy coordination work in 2024 and 2025 will become less tolerant in 2026 and 2027. The market correction is not hypothetical; it is already visible in hiring patterns at the organizations furthest along the adoption curve.

Transformation is harder than protection. It requires giving up comfortable visible work in exchange for less visible, more uncertain, higher-value work. It requires letting go of the identity that was built around producing the artifacts. It requires investing in the skills — systems thinking, stakeholder reading, difficult conversation, strategic framing — that were never formally taught in any PM or Scrum Master certification.

The tooling is the easy part. The rebalancing is the paper.

You can choose Path B. The judgment-heavy path has never been more valuable, and it is available to anyone willing to do the rebalancing work. But the work is real, and the window is shorter than it looks.

For practitioners and leaders.

For coordination practitioners

In the next 30 days

- 1 Audit your calendar honestly.** Two weeks of retrospective review. Categorize every hour as movement, judgment, or learning. Compute the ratio. If it's close to 70/30, you are where most of the market is, and you need to move.
- 2 Identify your two heaviest movement rituals.** For most Scrum Masters: status reporting and DoD verification. For most PMs: roadmap artifact production and stakeholder reporting. For most project managers: RAID log maintenance and schedule updates. Pick the top two and experiment with AI-assisted workflows that collapse them to a fraction of current time.
- 3 Reinvest the recovered time deliberately.** Do not let it fill with more movement work. Schedule judgment-layer activities — a weekly one-on-one you've been deferring, a strategic conversation with a stakeholder you've been avoiding, a cross-team alignment you've been meaning to drive.

In the next 90 days

- 1 Run the full ceremony audit.** Walk through all nine ceremonies from Section 3. For each, identify the movement work AI should absorb and the judgment work you should expand. Build the workflow changes into your actual practice.
- 2 Establish a learning slice.** Put two hours a week on your calendar, recurring, for deliberate skill building. Treat it as protected time. It will feel like a luxury; it is not.
- 3 Talk with your leader about the rebalancing.** Most leaders are not yet thinking about coordination roles this way. Being the practitioner who raises the conversation — with specifics — is both good for the organization and good for your own visibility.

For leaders of coordination teams

In the next 30 days

- 1 Audit your coordination layer.** Look at your PMs, SMs, project managers, and BAs as a portfolio. Which of them are on Path A — artifact-heavy — and which are on Path B — judgment-heavy? The answer is informative about your organization's current ratio and about who will thrive or struggle over the next 24 months.
- 2 Stop rewarding artifact production.** Performance reviews, project updates, and visibility rituals in most organizations implicitly reward volume of output over quality of judgment. Fix the incentives before the market does it for you.
- 3 Identify your Path B practitioners and invest in them.** They are among your most valuable people, and they are the hardest to replace. Make sure they know it. Make sure they have room to do judgment work.

In the next 12 months

- 1 Redesign your coordination layer deliberately.** Fewer people, higher seniority, rebalanced job descriptions, different measurement systems. Aim for a coordination layer where every practitioner spends 60–70%+ of their time on judgment work, supported by AI-augmented tooling underneath.
- 2 Invest in learning and transition paths for Path A practitioners who can make the move.** Many can. Some cannot. The leaders who handle this transition humanely — with clear expectations, real support, and honest feedback — retain their best people. The ones who don't will lose Path B practitioners first, because Path B practitioners notice organizational drift faster than anyone else.
- 3 Reshape hiring.** The PM, SM, and project manager hires that made sense in 2022 are not the same hires that make sense in 2026. Update the scorecards. Interview for judgment, for strategic thinking, for stakeholder navigation, for the capacity to use AI as a leverage multiplier rather than a task accelerator.

The longer arc. Every major shift in how organizations operate has rewarded practitioners who saw the shift early and painfully disrupted practitioners who did not. AI is one of those shifts, and the coordination layer is one of the places the disruption will land most visibly. The practitioners and leaders who engage with the rebalancing now — honestly, with specifics — will be in a substantially different position in 24 months than those who wait.

About Kollabio

Kollabio is a design and engineering firm that creates human-driven, AI-empowered solutions for federal agencies, GovCon primes, and commercial enterprises. We combine **Experience Thinking** — practical design thinking with the how-to included — with Agile delivery and continuous deployment to ship measurable value in multiple increments and with continuous feedback loops.

Our core competencies span strategy, experience design, product development, cloud architecture, and DevSecOps. We hold 8(a), HUBZone, CMMI DEV ML3, CMMI SVC ML3, ISO 9001:2015, and ISO 20000-1:2018 certifications, and are a GSA Contract Holder (MAS) as well as a holder on CMS DASH BPA and SBA EADS BPA contract vehicles.

This is Paper 2 in Kollabio's *Ground Truth* series. Paper 1 — *Rewriting the Build* — is available on our website. If you are navigating the shifts described in either paper inside your own organization, we welcome the conversation.

CONTACT

sales@kollabio.com
571-295-7209

OFFICE

23710 Schooler Plz Ste 275
Room #2058
Brambleton, VA 20148