



Whitepaper

Make your bed

The non-negotiable weekly refresh, and why cadence is the one thing your team can always control

FTTM best practices series.

How the weekly schedule refresh turns McRaven's bed-making lesson into the most controllable discipline on any development program — and how a three-day operating rhythm makes it produce a pull-in decision every week.

Prepared by

lateralworks
FTTM methodology

Date

July 2026
Methodology paper

Online

lateralworks.com
FTTM best practices series

Table of contents

Make your bed

Abstract	03
01 The smallest task, the biggest signal	04
02 Your project's made bed	06
03 Cadence beats content	08
04 The three-day operating rhythm	12
05 The critical path, recalculated weekly	15
06 Five elements, six design principles	18
07 Standing up the cadence	21
08 From beds to programs	24
A One week, one report: a worked example	26
References	34

Core thesis. The weekly schedule refresh is your project's made bed. It is the one event whose existence depends only on the team's own discipline. Never cancel it. Run on a fixed three-day rhythm, it proves every week that the team can show up, report honestly, update the schedule, review the critical path, and leave with a specific pull-in decision — and it turns the schedule from a record of the past into an instrument for changing the outcome.

Overview

Abstract

In May 2014, Admiral William H. McRaven told the University of Texas graduating class to start every day by making their bed. His point was not tidiness. It was agency: completing one small, controllable task creates momentum that carries through every task that follows [1]. This paper draws a direct line from that lesson to the practice at the heart of every program lateralworks coaches — the weekly schedule refresh.

Every complex development program operates inside a fog of uncontrollable variables: supplier delays, silicon errata, specification changes, resource constraints. Teams cannot control those events. They can control one thing absolutely — the cadence of their weekly schedule refresh. The refresh is a standing discipline, same day, same time, every week, in which every task owner reports an honest estimate against the schedule, the critical path is recalculated, and the core team decides on a specific pull-in action. The team walks out knowing exactly where the program stands.

The paper makes three arguments. First, cadence matters more than content: the refresh derives its power from being non-negotiable, and the behavioral science behind small wins, peer accountability, and habit formation explains why [2, 3, 4, 5]. Second, the refresh only earns its cost when it is operationalized — lateralworks codifies a three-day rhythm that moves from honest task-owner updates, through critical-path analysis by the program manager and product delivery lead, to a forty-minute pull-in meeting and an AI-assisted status report [6]. Third, consistent refresh execution builds compound returns in estimation accuracy, risk detection, and team confidence, visible week by week on the wiggglechart.

The argument rests on lateralworks' experience coaching semiconductor and systems development programs since 1988 — including a three-year, multi-hundred-million-dollar joint development program that missed exactly one refresh in three years — cross-referenced throughout against published research on execution cadence, operating rhythm, and schedule management [4, 7, 8, 9]. An appendix reproduces a complete AI-generated weekly status report from a live program, end to end.

01

Foundation

The smallest task, the biggest signal

In May 2014, retired U.S. Navy Admiral William H. McRaven delivered a commencement address at the University of Texas at Austin that has since been viewed more than ten million times [1, 10]. His opening lesson was disarmingly simple: make your bed every morning. In SEAL training, instructors inspected every recruit's bed before anything else. Corners squared, covers tight, pillow centered. A mundane act — and McRaven's point was not about tidiness. It was about agency.

Completing that first small, controllable task gave each trainee proof of accomplishment that carried through the rest of the day. It established a rhythm. It created evidence that discipline existed, and it gave people a reason to believe the next task was achievable. Accomplishing the first task, McRaven said, encouraged completing another, and then another.

The science under the anecdote

McRaven's lesson is a field observation, but the mechanism it describes is one of the most replicated findings in organizational psychology. Karl Weick argued in 1984 that people are paralyzed by problems framed at overwhelming scale, and that a strategy of “small wins” — concrete, complete outcomes of moderate importance — builds patterns that attract allies and lower the arousal that blocks clear thinking [2]. A slipping two-year program is exactly such an overwhelming problem; a single week's honest update is a small win.

Teresa Amabile and Steven Kramer reached a parallel conclusion from nearly 12,000 daily diary entries by 238 professionals on creative teams: of everything that distinguishes people's best working days, the single strongest factor is visible progress on meaningful work — even minor progress [3, 11]. Their “progress principle” is McRaven's cascade, measured: small, regular evidence of forward motion compounds into engagement, and engagement compounds into performance.

Charles Duhigg calls routines with this cascading property “keystone habits” — practices whose value lies less in their direct output than in the chain reactions they start [12]. Bed-making is his canonical household example. And the habit-formation literature explains why the discipline must be consistent to stick: Phillippa Lally's field study found that automaticity grows with repetition in a consistent context, reaching plateau after a median of 66 days — roughly the first quarter of running a new weekly discipline [5].

This paper draws a direct line from that body of evidence to a lateralworks best practice that sits at the heart of every program we coach: the weekly schedule refresh. The refresh is the program's keystone habit — the smallest task that carries the biggest signal.

02

The practice **Your project's made bed**

Every complex development program — semiconductor, systems, or software — operates inside a fog of uncontrollable variables. Silicon may fail at characterization. A supplier may miss a delivery window. A specification may change under the team's feet. Teams cannot control these events. But they can control one thing absolutely: the cadence of their weekly schedule refresh.

What the refresh is

The refresh is a standing discipline, same day, same time, every week, in which every task owner on or near the critical path provides a realistic assessment of remaining duration, the schedule is updated, and the critical path is recalculated. The output is binary and public: the predicted end date has pulled in, held, or slipped — and the team sees which, every seven days. Five elements are mandatory, and section 06 codifies them: fixed cadence, live update, critical path review, end-date tracking on a wigglechart, and backup ownership so the cadence is never broken [6].



Figure 1. The weekly refresh cycle. Five mandatory elements executed in sequence every week; backup ownership closes the loop so the cadence is never broken.

What the refresh is not

It is not a status call, and it is not a daily standup. A standup coordinates the work in flight today; a steering review allocates money and people quarter by quarter. The refresh does something neither of them does: it re-derives the program’s predicted end date from honest bottom-up estimates and forces a decision about how to move that date earlier. A meeting that ends with “we are three weeks behind” is a status meeting. A refresh ends with “here is the pull-in we chose on the critical path, here is the owner, here is the date” [6].

The distinction matters because the failure mode of schedule governance is almost never the absence of meetings. Programs drown in meetings while the master schedule quietly fossilizes. A schedule that is not refreshed weekly becomes a historical document — it describes what was planned, not what is real. The refresh transforms it into a navigation instrument the team recalibrates every week, the way lateralworks’ long pole method treats the critical path as a target to attack rather than a report to file [13].

03

Why it works

Cadence beats content

McRaven's bed-making ritual works because it is non-negotiable. It happens every morning regardless of what else is going on. The weekly refresh operates on the same principle, and its power comes from three properties: controllability, accountability, and biofeedback.

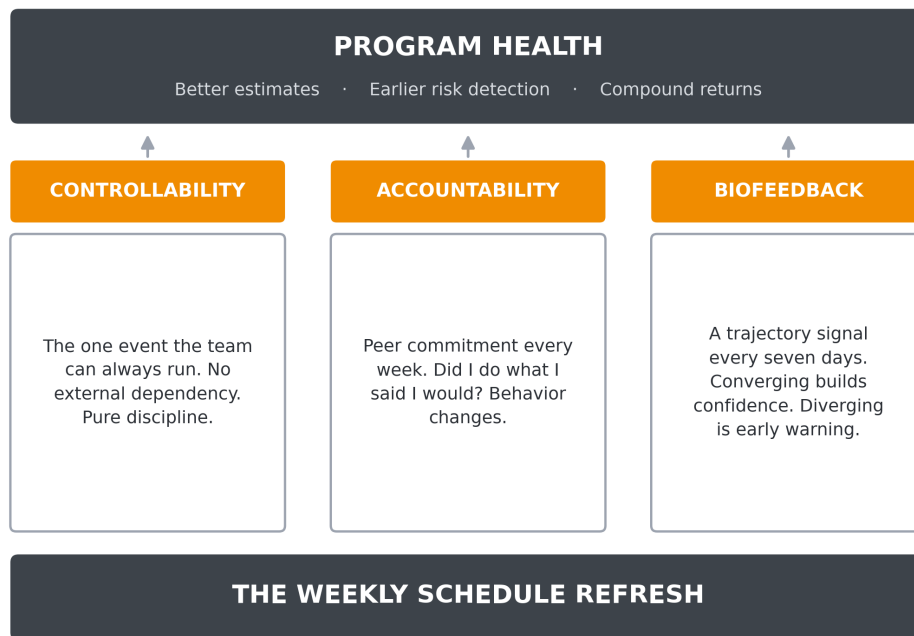


Figure 2. The three pillars of cadence. Controllability, accountability, and biofeedback compound over time into better estimates, earlier risk detection, and sustained program health.

Controllability

Every project has many uncontrollable factors. The team cannot prevent a supplier slip or a silicon bug. But running the refresh every week is entirely within the team’s control. It is the one scheduled event whose existence depends only on the team’s own discipline — which is precisely why it must never be canceled. When lateralworks coached a three-year, multi-hundred-million-dollar joint development program involving hundreds of people across multiple companies in the UK, the refresh was missed exactly once in three years: on Christmas Day [14]. That record was not an act of heroism. It was a made bed, three years running.

Accountability

When people know they will report progress to their peers every week, behavior changes. FranklinCovey’s 4 Disciplines of Execution reaches the same conclusion from the same direction: its fourth discipline installs a short weekly session, separate from every other meeting, in which each team member answers for last week’s commitments and makes next week’s — because commitments made to teammates, rather than only to a manager, become personal [4]. The refresh is that cadence of accountability applied to the master schedule. Each owner answers one question in public: did the work go as I said it would, and what is my honest date now?

Biofeedback

The refresh gives the team something most programs lack: a near-real-time signal of trajectory. After each update the critical path is recalculated and the predicted end date moves or holds. A team that watches its schedule pull in gains confidence and momentum — exactly the cascade McRaven described, and exactly

the progress principle in Amabile and Kramer's data [3]. A team that watches it slip gets an early warning measured in days rather than months. Early warning is the difference between a corrective action and a crisis.

Why weekly — not daily, not monthly

Cadence should match the tempo at which decisions must be made from new information. When General Stanley McChrystal's Joint Special Operations Task Force fought a networked enemy, its operations and intelligence brief ran daily, ninety minutes, thousands of people across seventy locations — because combat intelligence decayed in hours [7, 15]. A development program's schedule truth does not change hour to hour, but it changes far faster than a monthly review can catch: a month is long enough for a slipped task to compound through three dependent tasks and harden into the new critical path. A week is the interval at which schedule-grade information actually changes — long enough for real progress to measure, short enough that no slip is more than seven days old when the whole team sees it. Daily coordination belongs in the standup; weekly prediction belongs in the refresh; monthly and quarterly allocation belongs at the portfolio, where lateralworks runs the same refresh logic at lower frequency [16].

The lesson

Start with the small task

**“If you want to
change the world,
start off by
making your bed.”**

Admiral William H. McRaven
University of Texas at Austin commencement, May 17, 2014

04

Mechanics

The three-day operating rhythm

A principle without mechanics is a poster. lateralworks operationalizes the weekly refresh as a three-day rhythm that fits inside a normal workweek and ends, every week, with a pull-in decision [6]. Day 1 produces an update request. Day 2 turns task-owner responses into critical-path analysis and pull-in options. Day 3 runs a forty-minute core team meeting and issues the status report.

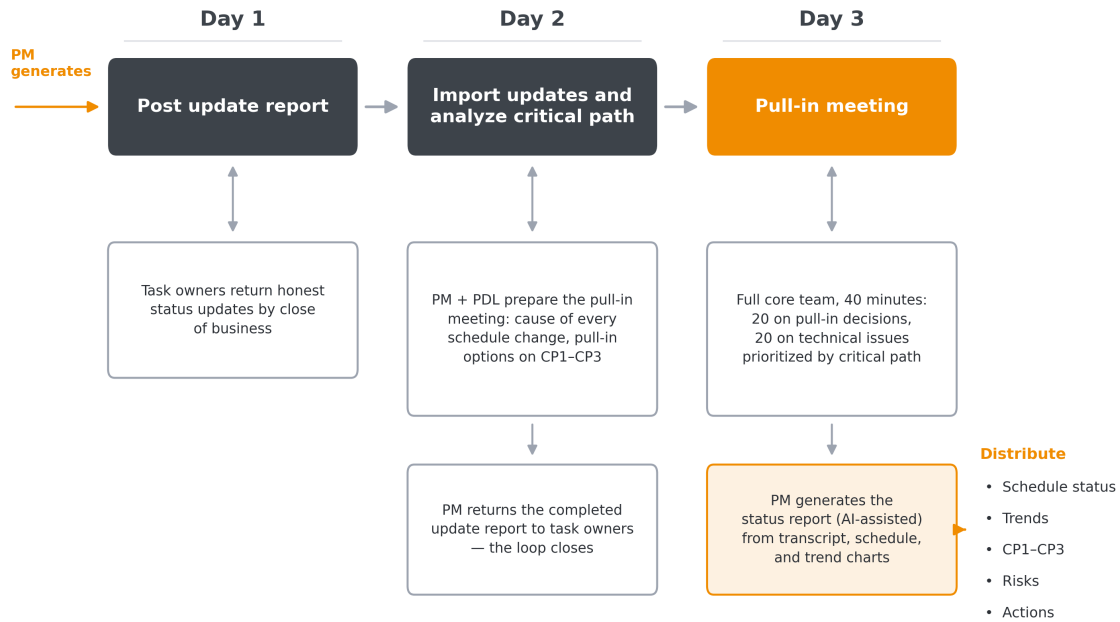


Figure 3. The three-day operating rhythm. The program manager’s outputs open and close the cycle; the pull-in meeting is where decisions are made and the schedule moves.

Day 1 — post the update report

The program manager (PM) generates and posts the update report to every task owner on the critical path and on any task scheduled to start, complete, or change state in the coming two weeks. Owners return their updates by close of business the same day. The report is intentionally short. For each task it asks three things: the current best-estimate finish date, the cause of any change from last week, and the top constraint blocking pull-in. It does not ask for percent-complete, because percent-complete invites negotiation. It asks for honest dates. The single-day window concentrates updates into one batch the PM can process on Day 2, instead of scattering them across chat threads and hallway conversations that never reach the master schedule.

Day 2 — import updates and analyze the critical path

The PM imports every returned update into the master schedule, then pairs with the product delivery lead (PDL) on critical-path analysis: which tasks now define the top three critical paths (CP1, CP2, CP3), what changed week over week, why, and where pull-in is possible on each path. The pairing is deliberate. The PM owns schedule mechanics and analytical discipline; the PDL owns engineering judgment — whether a pull-in option is actually feasible or a process step gates it. Neither role can prepare the pull-in meeting alone. Before the day ends, the PM sends the completed update report back to the task owners, so every owner sees their input reflected in the refreshed schedule. Visibility is what keeps the next round of updates honest.

Day 3 — the pull-in meeting

The full core team — the small, role-based team lateralworks puts at the center of every fast program [17] — meets for forty minutes. The first twenty go to pull-in: the PM and PDL walk the team through CP1 through

CP3, the cause of each change, and the candidate pull-in options identified on Day 2. The core team decides. Actions get named owners and dates. The second twenty minutes go to technical issues, prioritized by critical-path impact; issues that do not touch the top three paths are deferred to a separate technical forum. Immediately after the meeting, the PM generates the weekly status report. Generation is AI-assisted: the inputs are the meeting transcript, the refreshed schedule, and the wigglechart trend visuals, and the structure is fixed — schedule status, trends, CP1–3, risks, actions. One document, five sections, the same shape every week. Appendix A reproduces one in full.

Roles in the cadence

Role	Responsibility in the cadence
Program manager (PM)	Generates the update report (Day 1). Imports updates, analyzes the critical path, prepares pull-in options with the PDL, and closes the loop with task owners (Day 2). Runs the pull-in meeting and generates the AI-assisted status report (Day 3).
Product delivery lead (PDL)	Owns technical plausibility. Pairs with the PM on Day 2 to assess which pull-in options are engineering-viable. Leads the technical half of the pull-in meeting (Day 3).
Task owners	Return honest status updates by close of business on Day 1. Execute pull-in actions assigned on Day 3. See their updates reflected in the refreshed master schedule on Day 2.
Core team	Attends the forty-minute pull-in meeting. Makes pull-in decisions on CP1–3, resolves critical-path technical issues, and accepts the resulting actions and owners.

Weekly inputs and outputs

Day	Inputs	Outputs
Day 1	Prior-week master schedule; list of tasks in scope for update.	Update report published to task owners; owner responses returned by close of business.
Day 2	Task-owner responses; wigglechart trend data; prior-week critical-path analysis.	Refreshed master schedule; CP1–3 analysis with named causes; pull-in option slate; completed update report returned to task owners.
Day 3	Refreshed schedule; CP analysis; pull-in options; meeting transcript.	Pull-in decisions with owners and dates; updated risks and actions; AI-assisted status report distributed to stakeholders.

05

The instrument

The critical path, recalculated weekly

The refresh earns its forty minutes because of what it recalculates. The critical path method identifies the longest sequence of dependent activities from project start to finish — the one chain with no slack, the chain that sets the ship date [8, 18]. Whatever happens to that chain happens to the program. Everything else just watches.

The method is old enough to trust. Morgan Walker of DuPont and James Kelley of Remington Rand built critical-path scheduling between 1956 and 1959 to attack the cost of plant shutdowns, and DuPont cut shutdown costs by roughly a quarter with it [8]. What has changed since 1959 is not the math but the cadence at which organizations bother to rerun it. The U.S. Government Accountability Office’s schedule assessment guide makes “updating the schedule using actual progress and logic” and “confirming that the critical path is valid” two of its ten best practices for a reliable schedule — a schedule that is not updated with actual progress at regular intervals cannot tell the truth about the program it claims to describe [9].

The critical path is not static. It shifts as tasks complete faster or slower than planned, as scope changes, and as dependencies are re-sequenced. Near-critical paths — those within days of the critical chain — can silently become the new critical path if left unmonitored. lateralworks’ field data puts the danger zone at about ten days of float: paths that close sit one bad week from becoming the new long pole, which is why fast teams treat near-critical paths as critical and why the refresh reviews CP1 through CP3 every week, not CP1 alone [13, 19]. Weekly recalculation catches these shifts while corrective action is still cheap.

The wigglechart: the trend is the instrument

The projected end date is recorded at every refresh. Over time this produces the wigglechart — the weekly plot of predicted finish against target, and the single most information-dense artifact in the FTTM toolkit [13]. The line wiggles; that is the name and the point. An honestly updated schedule moves every week. Each slip shows, gets attacked in the room, and recovers within weeks. A line that never moves is a schedule that has stopped speaking — a flat line is not stability, it is silence.

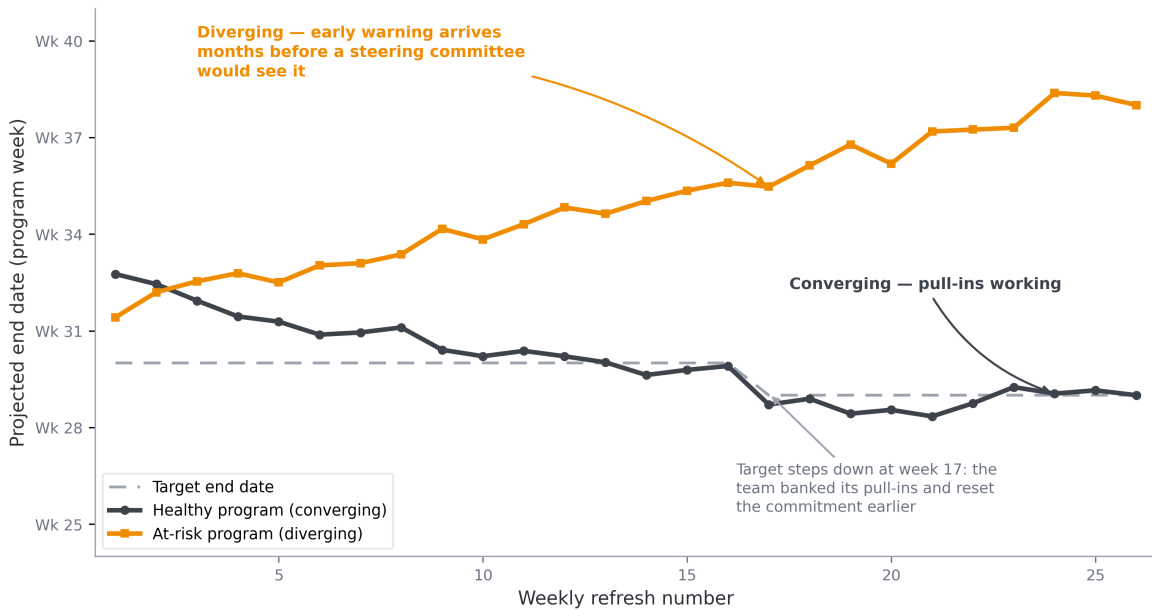


Figure 4. Wigglechart: end-date trend over 26 weekly refreshes. The slate line converges as pull-ins land and banks the surplus by stepping the target down at week 17. The orange line diverges — and says so months before a quarterly review would.

Read this way, the wigglechart is the early warning system every late project wishes it had bought in time. The comparison with a monthly reporting cycle is stark: a slip that lands on the wigglechart within seven days can hide inside a monthly cycle for four to eight weeks, filtered through status-slide optimism at every hop.



Figure 5. Feedback-loop compression. The same week-2 slip stays dark for weeks under a monthly review cycle; under a weekly refresh it is public within seven days, while corrective action is cheap.

06

Codification

Five elements, six design principles

Based on coaching more than two hundred development programs since 1988, lateralworks codifies the weekly refresh as five mandatory elements carried by six design principles. The elements are what a team does; the principles are why the mechanics hold up under pressure [6, 20].

The five mandatory elements

Fixed cadence. Same day, same time, every week. No exceptions. If the meeting owner is unavailable, a trained backup runs it. This is the make-your-bed discipline: the meeting happens regardless.

Live schedule update. The schedule is updated during the cycle, not reconstructed after it. Each task owner provides a current honest assessment of remaining duration, and owners see the impact of their updates in the refreshed schedule the next day.

Critical path review. After updates land, the team reviews the recalculated critical path. Has it changed? Are new tasks critical? Are near-critical paths closing in? The team agrees on pull-in actions where possible.

End-date tracking. The projected end date is recorded at every refresh and plotted on the wiggglechart. A converging trend signals a healthy program; a diverging or oscillating trend signals systemic issues that need intervention.

Backup ownership. Every refresh facilitator has a designated, trained backup who runs the cycle without hesitation when the primary is sick or traveling. The cadence is never broken.

The six design principles

SME ownership of estimates is preserved. Task owners give honest best-estimate dates every week without being pushed to match a target. The gap between estimate and target stays visible in the schedule — knowing the gap is how you close it. The moment owners learn that estimates get negotiated to a target, the estimates stop being honest and the gap stops being visible.

The critical path drives the agenda. Most schedule meetings allocate equal time to every task. This one does not. CP1 through CP3 get the airtime; everything else is deferred. That discipline is why forty minutes is enough. The core team's attention is the scarcest resource in the program, and the cadence spends it only where it can change the outcome.

Pull-in, not status. Status meetings end with “we are N weeks behind.” Pull-in meetings end with a chosen option, an owner, and a date. The cadence is named for what it produces: a week without a pull-in decision is a week the cadence did not run, even if the meeting took place on time. The same rule powers the long pole method's refresh planning: attempt a pull-in every week, and bank the surplus so later slips land on the bank instead of the date [13].

AI-assisted reporting keeps the load sustainable. The weekly status report is assembled from the meeting transcript, the refreshed schedule, and the trend charts, leaving the PM's time for analysis rather than typing. Without that automation, a weekly cadence collapses under its own administrative weight within a month. With it, the marginal cost of the weekly report approaches zero while the format stays identical week over week — Appendix A shows the result.

The loop closes on the task owners. Every owner sees their Day 1 update reflected in the Day 2 refreshed schedule. That is why the next Day 1 update is honest. Cadences that take inputs and never report back break the trust that keeps status honest, and the inputs degrade into what people think leadership wants to hear.

Fixed day, fixed roles, fixed artifacts. The same day of the week, the same meeting length, the same report structure, the same distribution list. Variation costs attention; consistency buys it back. By Monday the core team should know exactly what the cadence will produce by Friday. This is the “consistent context” habit research identifies as the precondition for automaticity [5].

Normal organization vs. best organization

Normal organization	Best organization
Schedule reviewed monthly or quarterly; slips surface on status slides weeks after they occur.	Schedule refreshed weekly; every slip is public within seven days, raw, on the wigglechart.
Meetings end with a status: “we are N weeks behind.”	Meetings end with a decision: a pull-in option, an owner, and a date on CP1–3.
Estimates negotiated toward the target; the gap disappears from view and reappears at the end.	Honest bottom-up estimates; the gap stays visible and gets attacked weekly.
Refresh canceled for holidays, travel, crunch weeks — cadence breaks under pressure.	Trained backup runs the refresh; one miss in three years, on Christmas Day.
Reporting consumes the PM; analysis happens if time remains.	AI-assisted report generated from transcript, schedule, and trends; PM time goes to analysis.

07

Implementation

Standing up the cadence

The cadence assumes four things exist before week one: a core team that is formed, chartered, and meeting regularly [17]; a master schedule in a tool that supports weekly updates and critical-path analysis; a wiggglechart or equivalent trend visual that records milestone movement week over week; and a PM and a PDL — distinct individuals, each with time to do the Day 2 work well. If any one is missing, put it in place before starting. The cadence consumes standing artifacts; it does not recreate them.

Standing artifacts

Four artifacts support the cadence from week one: the master schedule, the wiggglechart, the risk register (feeding the risks section of the status report), and the action log (feeding the actions section). These are maintained continuously, not assembled weekly. The cadence updates them and consumes them.

The first four weeks

The first two weeks feel thin. Critical-path analysis is still calibrating, pull-in options are shallow, and task owners are learning what an honest update looks like. That is normal. By week three the updates stabilize. By week four the wiggglechart begins to show the trend: either the gap is closing, or it is not. Either answer is valuable — if the cadence is surfacing that the team cannot close the gap, it is doing its job. It has turned an unnamed drift into a visible management problem.




Figure 6. The calibration curve. Estimate quality and refresh maturity build asymptotically over the first month — the same shape habit research measures for automaticity under consistent repetition.

Habit research offers one more encouragement and one warning. Lally’s participants who missed a single repetition did not reset the curve — one broken week does not destroy the discipline [5]. But automaticity only built at all when the behavior repeated in a consistent context. Move the meeting around the calendar and the habit never forms. Fixed cadence is not a preference; it is the mechanism.

Failure modes to watch

The cadence fails in recognizable ways. Day 1 updates that trickle in after Day 2 mean the deadline is not being enforced. A pull-in meeting that spends its forty minutes on tasks outside CP1–3 means the PM–PDL preparation did not hold the agenda. A status report that never reaches the task owners is a broken output loop. And a core team that stops making pull-in decisions, week after week, points to a leadership problem the cadence cannot fix on its own — the instrument is telling the truth; someone has to act on it.

The test. Every week, one question answers whether the cadence is working: did the program make a specific pull-in decision on CP1, CP2, or CP3 this week? If yes, the cadence is doing its job. If the answer is no for several weeks running, the cadence is running without producing — and the root cause is upstream of the cadence.

08

Conclusion

From beds to programs

McRaven's lesson resonates because it reframes discipline as a source of confidence rather than constraint. Making your bed does not change the world. But proving to yourself, every morning, that you can complete one task to standard — that you control at least one thing — creates the foundation for everything else [1, 10].

The weekly refresh is your project's made bed. It proves that the team can show up, report honestly, update the schedule, review the critical path, and walk away with a pull-in decision. Every week. Without exception.

The discipline builds compound returns. Each week the team practices accountability, it gets better at estimating. Each week it reviews the critical path, it gets sharper at spotting risk. Each week it sees the trend line, it develops an intuitive sense of program health that no monthly steering committee can replicate. And the logic scales: the same refresh discipline that keeps one program honest runs at monthly cadence across the portfolio, where lateralworks applies it to project starts, rankings, and resource gates [16]. Cadence is not one practice among the ten that fast teams share [20] — it is the practice that keeps the other nine honest.

Start small. Fix the day and the time. Post the first update report on a Day 1. Run the first forty-minute meeting on a Day 3, even if the analysis is thin. Then do it again next week, and the week after. The wigglechart will take care of the rest.

Key principle. Never cancel the refresh. It is the one thing you can control. Make your bed.

A

Appendix

One week, one report: a worked example

What follows is a complete, unedited weekly status report from a live semiconductor development program (identifying details generalized). It is the Day 3 output of the cadence described in section 04 — generated by AI skills within minutes of the pull-in meeting, from four inputs the cadence maintains anyway: the refreshed master schedule, the meeting transcript, the action log, and the risk register.

Because the generation is skill-driven, the report has the same five sections in the same order every week: status summary, schedule and trends, this week's tasks, actions, risks. Readers stop hunting and start comparing — week-over-week deltas jump out because the frame never moves. Note what the report contains that a hand-built status slide rarely does: named causes for every schedule change, the top three critical paths with float, a first material pull-in (23 days on the qualification milestone) traced to the specific replan that produced it, and seventeen open actions with owners and dates.

Weekly schedule refresh

Weekly program status • 22 April 2026

1. Status summary

First significant pull-in of the program — 23 days on the qualification milestone. The new quality lead completed a detailed replan of the qualification phase, cutting execution from six months to four months (240d → 150d) and qualification report writing from 60 days to 30 days. In parallel, a **shared-risk qualification path** at 168-hour HTOL has been added to enable the lead customer to start application qualification in parallel with full product qualification. Net effect: WS: Product qualified predicted pulls in from 30-Apr-28 to 7-Apr-28, closing the gap to target from 152 to 129 days. WS: Samples available shows no change this week — still predicted 3-Sep-27 (65-day gap). FS wafers have not been re-baselined yet despite the foundry partner running ahead; the program manager plans to break down the 117-day foundry task into process stages over the coming weeks and expects a further 1–2 week pull-in once the data is ingested.

Milestone	Target	Predicted	Gap	Change
WS: Samples available	30-Jun-27	3-Sep-27	(65)	No change
Product qualified	30-Nov-27	7-Apr-28	(129)	+ 23 days (pull-in)

Key wins this week: Foundry-partner management meeting held — four company attendees including the program manager had a productive, candid exchange; the foundry confirmed the 3-day/mask-layer assumption, shared their internal dZ waterfall chart, and set up a bi-weekly cadence. Actual wafer progress is tracking slightly ahead of the foundry's own schedule. **Second critical path** (TC1 wafer starts) was pulled in 11 days on the back of the foundry's confirmation — now at 35 days from CP (was 24 days). OSAT-partner tester/prober spec progress continues but the sourcing lead encountered a commercial-terms snag on the term sheet; he is splitting engineering and commercial tracks to unblock engineering access immediately. **Customer engagement escalated** — the customer contact responded, meeting with the customer COO secured for April 28; the test engineering lead is finalising functional test spec this week with a customer-ready version next week. The new quality lead has formally taken ownership of all qualification tasks (Task 324+), closing the handover action from last week.




Figure 1: Program timeline

2. Schedule and trends

Critical path. The critical path is structurally unchanged through to WS samples but the tail has been compressed. Sequence: FS: Wafer starts to ALD-DEP (117d, process lead, 0 float) → ALD vendor processing (30d) → Reinsertion (63d) → First-silicon wafers out (11-Nov-26) → Validate sort (30d, validation engineer) → Ship & assemble (7d) → Final test (30d) → Samples available (17-Jan-27) → Validate design (package, 45d, design lead) → WS redesign (30d) → Final masks → WS samples available (3-Sep-27). Downstream, the **qualification tail has been shortened**: Execute Qualification now 150d (was 240d), Complete Qual report now 30d (was 60d), Qualification Done now 7-Apr-28 (was 30-Apr-28). The phase has also been renamed *Qualification (3rd respin)* to reflect the planning basis. Critical attention remains on the 117-day foundry fab window and the August 10 ALD-DEP hold date; decomposition of the foundry task into FEOL / gate / contact / metal stages is the next schedule improvement.




Figure 2: Critical path (extracted from program schedule)

Wiggle chart. The qualification milestone shows the **first material pull-in since tracking began** — predicted finish moves from 30-Apr-28 to 7-Apr-28 (+23 days, green trend arrow). The risk indicator on qualification has flipped to pull-in. WS: Samples available shows zero change this week: trend holds at 3-Sep-27 with the 65-day gap intact, because the FS critical chain has not been re-based against the latest foundry data yet. Cause of change recorded: *"Pull-in due to reassessing qualification timelines with new quality lead."* This is a directionally positive week but the structural 65-day gap on WS samples remains the governing constraint.



Figure 3: Wigglechart dashboard (extracted from refresh tracker)

Critical path analysis. Significant reshuffling of the near-critical paths this week. **CP2 and CP3 have swapped:** last week's CP2 (TC1: Wafer starts to ALD-DEP, process lead) was pulled in 11 days on the back of foundry confirmation and is now CP3 at 35 days from CP. The new CP2 (31 days from CP) is a test-hardware-lead / COO cluster: Sort 1 tester/prober spec from the OSAT partner (2d), Sort 1 probe-card design (14d), and FS procurement approval for package testing infrastructure (7d, COO). CP4 (37 days from CP) shifted by 90 days due to a newly-discovered package-test-spec dependency but retained its margin; it now runs through the test engineering lead's Sort 1 functional test spec (46d) and the offshore package test lead's package-test-fixture design (45d). The team agreed to continue monitoring CP2 through CP4 weekly and to watch whether the test-hardware-led cluster creeps closer to zero float as probe-card design kicks off.

Other schedule updates

- **Qualification phase replanned by quality lead.** Execute Qualification compressed from 240d to 150d (6 months → 4 months); Complete Qual report from 60d to 30d. Qualification start now gated on Design validated rather than Samples available, which the team flagged as the technically correct dependency. Phase renamed *Qualification (3rd respin)*.
- **Shared-risk qualification path added** at 168-hour HTOL (vs. standard 1000-hour), enabling the lead customer to begin application qualification in parallel with product qualification — saves the customer approximately 2 months. Activation-energy and acceleration-factor calculations to be confirmed by the quality lead.
- **Foundry management meeting held.** Open communication with the foundry partner, bi-weekly cadence set, foundry confirmed 3-day/mask-layer commitment, foundry shared their internal dZ waterfall chart. Company tracking aligned with foundry tracking. The process lead is updating the mask-level crawl chart 3× per week and will share via meeting minutes.
- **OSAT term sheet — commercial snag.** The OSAT partner wants commercial terms agreed before releasing technical contacts. The sourcing lead is restructuring into two parallel tracks: engineering pricing agreement first (to unblock engineering access), commercial / production terms in parallel. Estimate remains ~2 weeks. Successor linkage of OSAT tasks to the signed term sheet still outstanding — action assigned to the program manager.
- **ALD vendor engagement meeting** scheduled for May 5 to align on ALD demo window. Original vendor readiness was 23-Jul which is earlier than wafer arrival; current discussions exploring combined FS+TC1 ALD shipment. Vendor lab is not high-volume and protocol maturity will be checked on-site; live demo required.
- **Customer engagement escalating.** Meeting secured with the customer COO on April 28. The customer requires “organisational maturity” evidence — product spec, detailed test spec, and qualification spec. The test engineering lead is finalising functional test spec this week; customer-condensed version to follow next week. Customer milestones not yet carried in the schedule — deliberate, pending customer-engagement maturity.
- **Design validation plan debate.** The customer engagement lead and the sourcing lead raised that the 45-day design validation block has no formal plan document, making duration a low-confidence estimate. The test engineering lead confirmed a bring-up / validation test list exists but is not formalised as a plan. Action: test engineering lead to align with the design lead and share; team to decide next week whether a formal validation plan task is added.
- **Package test spec dependency added.** New predecessor chain identified in-session — the offshore package test lead's package-test-fixture design feeds the package test spec, which feeds the test hardware lead's package test program development. Critical path 4 absorbed the shift without losing float.
- **Foundry actual wafers ahead of schedule.** Actual wafer progress (per foundry's dZ waterfall) is slightly ahead of their own schedule at early mask levels. Cautious optimism — 3-day/mask tracking established as the crawl-chart baseline.
- **Second lot tracking.** The foundry's second lot (staggered by two weeks) continues; primary schedule tracks lot 1 only; team tracking both separately.

3. This week's tasks

The **process lead's** FS: Wafer starts to ALD-DEP (CP 1, 0 days float, 117 days) remains the program driver — the foundry lot is running and mask-level tracking is now in place. The **test hardware lead** now owns the new CP 2 cluster (31 days from CP): obtain OSAT tester and prober spec (2d, overdue), issue Sort 1 probe-card PO (quote in hand from the test-card vendor), then design Sort 1 probe card (14d). The **COO's** FS procurement approval for package testing infrastructure (7d) sits on CP 2 and must precede the package-test-fixture PO with its 150-day lead. TC1: Wafer starts to ALD-DEP (**process lead**, CP 3, 35 days from CP, 125d duration) is the second wafer path to watch. The **test engineering lead** is driving Sort 1 functional test spec (CP 4, 46d, 37 days from CP) — target finish 24-Apr — and will then produce customer versions for the lead customer. The **offshore package test lead** owns package-test-fixture design (CP 4, 45d). The **tool vendor liaison** has TC1: Secure agreement with the ALD vendor (CP 10, 5d) — this is the procurement gate for the TC1 ALD slot. The **sourcing lead** is running OSAT term-sheet negotiation (CP 9, 22d) with restructured engineering-first approach. The **validation engineer** hit SPI2: Sort 1 (CP 8, 9d) this week.

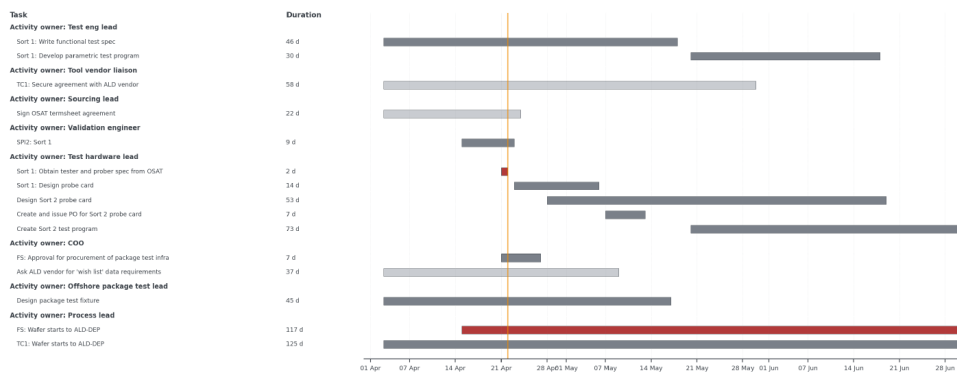


Figure 4: Focus report — work by owner (extracted from program schedule)

4. Actions

Prior actions closed: Transfer qualification tasks (Task 340+) to the new quality lead (program manager / test engineering lead / quality lead, due Apr-26) — closed: quality lead now owns all qualification tasks in the schedule and has replanned the phase. Break down the foundry's 117-day fab task and prepare for the April 27 foundry management meeting (program manager / process lead, due 27-Apr) — meeting held, closed; follow-up decomposition into FEOL / gate / contact / metal stages rolled into the open "Confirm foundry estimate" action. Set up foundry lot-tracking (process lead, due Apr-26) — closed: bi-weekly cadence and 3x/week crawl-chart tracking now established.

New / open actions:

#	Owner	Action	Due	Source
1	Validation engineer	Discuss mitigation options for ALD deposition (WS: Samples available). Carried over multiple weeks — needs resolution or formal closure.	3-Apr-26 OVERDUE	Action log
2	COO	Push the foundry to process faster — management meeting held; continue pressure on lot prioritisation through new bi-weekly cadence.	27-Apr-26	Action log
3	Program manager	Follow up with the sourcing lead on signing the OSAT termsheet agreement to assign successor links in the schedule (signed term sheet is a predecessor to multiple OSAT tasks).	29-Apr-26	Action log
4	Program manager	Confirm foundry estimate of earlier delivery for FS: Wafer starts to ALD-DEP. Ingest mask-level data and break the 117-day task into FEOL / gate / contact / metal stages.	5-Jun-26	Action log
5	Process lead	Align with the foundry liaison on argumentation for the foundry to shorten Reinsertion time (FS: Reinsertion to wafer out) — leverage the sourcing lead's ALD contamination-protocol data.	5-Jun-26	Action log
6	Sourcing lead	Execute split OSAT strategy — send engineering pricing list, get OSAT to release engineers immediately, pursue commercial / production agreement in parallel. Target completion ~2 weeks.	6-May-26	Schedule
7	Test engineering lead	Complete Sort 1 functional test spec (target 24-Apr) and produce customer-ready condensed version (aim for next week). Have the sourcing lead sanitise before sending.	1-May-26	Schedule
8	Test engineering lead	Plan out all test specification milestones (parametric, functional, package test, qualification spec) and share with the COO so dates can be communicated to the customer.	29-Apr-26	Schedule
9	Test eng lead / Design lead	Align on and share design validation plan / bring-up test list. Team to decide next week whether a formal validation plan task should be added to the schedule.	29-Apr-26	Schedule
10	Quality lead	Confirm activation-energy and acceleration-factor calculations for the 168-hour shared-risk HTOL qualification path; document customer-facing assumptions.	6-May-26	Schedule
11	Test hardware lead	Obtain tester and prober spec from the OSAT partner (overdue). Issue probe card design PO upon receipt of spec (quote in hand from the test-card vendor). Kick off Sort 1 probe-card design (14d).	29-Apr-26	Schedule
12	Test hardware lead / COO	Secure COO procurement approval (7d) for package testing infrastructure (150-day lead). Place PO immediately after approval. Maintain OSAT test-time rental as fallback.	15-May-26	Schedule
13	Process lead	ALD vendor meeting 5-May — prepare for ALD demo alignment; push the vendor for cut-off date on TC1 ALD window re-slotting; evaluate combined FS+TC1 ALD shipment option.	5-May-26	Schedule
14	Process lead / Program manager	Decide on alternative ALD partner vs. primary vendor lab. Meeting deferred from previous Friday.	25-Apr-26	Schedule
15	Test engineering lead	Align with offshore package test lead on package-test-fixture spec and test-program development — these activities need to start now given long lead times.	29-Apr-26	Schedule
16	Program manager	Propose an async update mechanism for the offshore package test lead given timezone conflict with the weekly refresh meeting. Ensure their input is captured weekly.	29-Apr-26	Schedule
17	COO / Customer engagement lead	Following April 28 customer COO meeting, determine when customer milestones can be formally added to the schedule (currently held out).	8-May-26	Schedule

Appendix A, page 6 of 7. AI-generated weekly status report — example deliverable.

5. Risks

Risk	Description	Severity
Structural schedule gap	WS: Samples target 30-Jun-27 vs. predicted 3-Sep-27 — 65-day gap (unchanged this week). Product qualified target 30-Nov-27 vs. predicted 7-Apr-28 — 129-day gap (improved by 23 days via the quality lead's qualification replan and 168h HTOL shared-risk path). Structural gap on WS samples remains the governing constraint — further foundry-task decomposition expected to yield 1–2 weeks more, but closing the 65-day gap still requires replan or parallelism change.	Critical
Foundry fab cycle time	Now materialising but trending better . FS lot running since April 15; foundry's actual wafer progress tracking slightly ahead of their own schedule at early mask levels; 3-day/mask commitment confirmed; dZ waterfall visibility gained; crawl-chart tracking 3x/week. Bi-weekly cadence established. Remaining concerns: 117d is still a single monolithic task pending decomposition; reinsertion 63d and PM windows remain uncertain; FS wafers have not been re-based against newer foundry data yet.	Critical
Test hardware lead times	Unchanged this week. Probe-card PCB 70d (was 42d); Package test infrastructure 150-day lead time; approval + PO + lead = ~167 days. Vendor still unresponsive. Probe-card PO unblocked by quote in hand but design gated on OSAT tester/prober spec. OSAT test-time rental remains fallback but cannot support design validation at speed.	Critical
OSAT dependency chain	Worsened slightly — the OSAT partner pushed back on releasing engineers before commercial terms agreed. The sourcing lead has restructured into parallel engineering + commercial tracks to unblock tester/prober spec delivery. NDA signed; PO quote in hand; tester/prober spec still pending. Successor linkage of OSAT tasks to the signed term sheet still outstanding (program manager follow-up).	Critical
Owner concentration	The process lead now owns CP1 (FS lot, 0d) and CP3 (TC1 wafer starts, 35d). The test hardware lead owns the new CP2 cluster (Sort 1 tester spec, probe card) and CP6 (Sort 2 tester). Load on these two individuals is a systemic risk — any absence or competing priority hits multiple paths. No deputy or backup author assigned yet for the most critical individual tasks.	High
Design validation scope	New this week. 45-day design validation block has no formal validation plan document; the customer engagement lead and the sourcing lead flagged this as a low-confidence duration and open-ended activity. The test engineering lead confirmed an informal bring-up / validation test list exists. Risk: validation drifts open-loop once started, amplifying downstream WS and qualification variance. Test eng lead and design lead to align and share; formal task TBD next week.	High
Customer alignment / spec readiness	New this week. April 28 meeting with the customer COO raises customer commitment; the customer requires product spec, detailed test spec, and qualification spec as evidence of organisational maturity. The test engineering lead is completing functional test spec this week; customer-ready condensed versions to follow. Current schedule does not carry customer milestones — could cause surprises as the customer's own qualification planning firms up.	High
Foundry escalation approach	Materialised and trending positive. CEO / management meeting held; bi-weekly cadence set; relationship preserved. Continuing risk: if wafer progress stalls, the escalation lever is now partially spent. Maintaining face with foundry management remains culturally important.	High
ALD vendor capacity & ALD window	ALD vendor meeting May 5. Original demo readiness 23-Jul is earlier than wafer arrival; combined FS+TC1 ALD shipment under discussion. Vendor lab is not high-volume; contamination-protocol maturity unverified; live demo and protocol audit required. Reinsertion contamination concerns (historical nickel data) add pressure.	High
Near-critical paths reshuffling	CP2/CP3 swapped this week. CP4 absorbed a 90-day shift without losing float. Schedule is still plastic and sensitive to small duration changes. Need to watch whether the test-hardware-led CP2 cluster closes on zero float as probe-card design kicks off.	Medium
Procurement gaps	Substantial progress on procurement breakdown completed over last two weeks. Remaining: cash-flow forecasting, backup OSAT qualification, and location-tagging of all schedule tasks (program manager).	Medium
BIST design issue	Potential built-in self-test problem flagged previously by the COO and sourcing lead. Still under investigation; no schedule impact yet.	Medium
Offshore communication	Time-zone conflict — early hours offshore for the weekly refresh meeting. No update received in the report this week. Tuesday sync with the test hardware lead continues but async mechanism needs improvement. Substrate design has not yet started despite NDA being signed.	Medium

Sources

References

- [1] McRaven, W. H. Commencement address, University of Texas at Austin, May 17, 2014. Transcript at news.utexas.edu; video at <https://www.youtube.com/watch?v=pxBQLFLei70>
- [2] Weick, K. E. "Small Wins: Redefining the Scale of Social Problems." *American Psychologist*, vol. 39, no. 1, 1984, pp. 40–49.
- [3] Amabile, T., and Kramer, S. *The Progress Principle: Using Small Wins to Ignite Joy, Engagement, and Creativity at Work*. Harvard Business Review Press, 2011.
- [4] McChesney, C., Covey, S., and Huling, J. *The 4 Disciplines of Execution: Achieving Your Wildly Important Goals*. Free Press, 2012. See Discipline 4: Create a Cadence of Accountability.
- [5] Lally, P., van Jaarsveld, C. H. M., Potts, H. W. W., and Wardle, J. "How Are Habits Formed: Modelling Habit Formation in the Real World." *European Journal of Social Psychology*, vol. 40, no. 6, 2010, pp. 998–1009.
- [6] lateralworks. "Weekly schedule refresh cadence: a 3-day operating rhythm." FTTM methodology brief, April 2026.
- [7] McChrystal, S., Collins, T., Silverman, D., and Fussell, C. *Team of Teams: New Rules of Engagement for a Complex World*. Portfolio/Penguin, 2015.
- [8] Kelley, J. E., and Walker, M. R. "Critical-Path Planning and Scheduling." *Proceedings of the Eastern Joint IRE-AIEE-ACM Computer Conference*, December 1959, pp. 160–173.
- [9] U.S. Government Accountability Office. *Schedule Assessment Guide: Best Practices for Project Schedules*. GAO-16-89G, December 2015. <https://www.gao.gov/products/gao-16-89g>
- [10] McRaven, W. H. *Make Your Bed: Little Things That Can Change Your Life... And Maybe the World*. Grand Central Publishing, 2017.
- [11] Amabile, T., and Kramer, S. "The Power of Small Wins." *Harvard Business Review*, May 2011. <https://hbr.org/2011/05/the-power-of-small-wins>
- [12] Duhigg, C. *The Power of Habit: Why We Do What We Do in Life and Business*. Random House, 2012.
- [13] lateralworks. "The long pole: find the path that sets the end date." FTTM methodology white paper, July 2026. lateralworks.com/papers.
- [14] lateralworks. Internal assessment database. Engagement data across client programs, 1988–2026.
- [15] Fussell, C. *One Mission: How Leaders Build a Team of Teams*. Portfolio/Penguin, 2017.
- [16] lateralworks. "Starts control: balancing the portfolio for speed." FTTM best practices white paper, July 2026. lateralworks.com/papers.
- [17] lateralworks. "The FTTM core team: roles, not functions." Team series white paper, 2026. lateralworks.com/papers.
- [18] Project Management Institute. "Critical Path Method Calculations — Project Schedule Terminology." [PMI.org](https://www.pmi.org).
- [19] Mitchell, N. "The critical path you're not watching." lateralworks briefing paper, May 2026. lateralworks.com/papers/the-critical-path-youre-not-watching.pdf
- [20] lateralworks. "10 best practices of fast teams." FTTM best practices series. lateralworks.com/papers/10-best-practices-of-fast-teams.pdf