



White Paper

Fungible Capacity

Why standing teams rebuild the old constraint.
The pool model for the AI-native organization.

Alexander S. Petty

SINGULARICS | singularics.ai

© 2026 SINGULARICS. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of SINGULARICS.

For permissions, bulk copies, or licensing inquiries:

hello@singularics.ai

First published August 2026.

singularics.ai

Contents

1	The Pool	4
2	Plan	6
3	Commit	7
3.1	Commitment, in practice	7
4	Release	8
4.1	Make release explicit	8
5	How an Effort Ends	9
5.1	The Decision Owner proposes closure	9
5.2	The Flow Council confirms	9
5.3	The one-week stall rule	9
6	Fluid Formation	11
6.1	Three triggers for reformation	11
6.2	Non-triggers	12
7	Standing Teams Rebuild the Old Constraint	13
7.1	The mechanism of silent failure	13
7.2	The correction	14
8	Failure Modes	15
8.1	Capacity never released	15
8.2	Trains reassemble	15
9	How the Pool Model Holds Together	16

Abstract

When delivery cost falls toward zero, organizations discover that their standing teams were never really about teamwork. They were about allocating scarce delivery capacity to systems. Remove the scarcity and the allocation becomes the constraint. This paper describes the pool model: one pool of people and agent fleets, capacity planned weekly and released the moment an outcome lands. It covers how the Flow Council plans, commits, and releases capacity, how an Effort ends, why fluid formation is narrower than the phrase suggests, and the silent failure that rebuilds permanent teams within two quarters if release is skipped.

The Pool

Ask any portfolio leader whether they can move capacity to their highest-value work this week. In most organizations the honest answer is no, because every person is assigned to a standing team and every team is assigned to a system. The capacity exists on paper. It does not move in practice. We have seen portfolios where 100% of delivery capacity was allocated, 40% of active initiatives had stalled, and the leadership team still described the problem as a “resource shortage.” The shortage was not people. It was the inability to redeploy the people they had.

The Flow Council holds one pool of capacity for the portfolio. The pool contains both people and agent fleets. Capacity is planned and released. It is never allocated to a standing team.

Nobody owns a system. That absence is the design choice that makes the rest of the model work. An organization that restores system ownership has lost the ability to move capacity, whatever the operating model says.

Capacity is planned and released. It is never allocated to a standing team. Nobody owns a system.

Traditional portfolio management allocates people to teams and teams to systems, then treats the resulting structure as a fact of organizational life. The pool model treats it as an artifact of the old constraint. When execution was scarce, you had to park capacity near the work it would eventually do, because moving it was

expensive and slow. That parking arrangement hardened into team boundaries, system ownership, and the assumption that a person belongs to a team rather than to a portfolio.

When execution is abundant and the constraint is human judgment, the parking arrangement becomes the bottleneck. A standing team sitting on a system that does not need attention this week is capacity the portfolio cannot use, and the portfolio feels it as a shortage even though the headcount has not changed.

The pool dissolves the parking arrangement. People and fleet capacity sit in one place, visible to one group, and flow to wherever the portfolio's highest-value outcome needs them.




Figure 1: The capacity cycle. Capacity is planned against the Now horizon, committed to an Effort, used during execution, and released back to the pool when the outcome lands. The cycle repeats weekly. Skipping the release step is how standing teams reassemble silently.

Plan

Each week the Flow Council sizes the human and agent capacity available against the Now horizon. The sizing is a subtraction, not an estimate. Start with headcount and subtract what is not available. What remains is capacity.

- **Leave.** Actual leave for the coming week, confirmed not forecasted. A person on leave for three of five days contributes 40% of a person-week, not 100%.
- **Standing obligations.** The recurring business duties that do not pause because someone joined an Effort. For most business-side staff this is 20 to 40% of their week. For Fleet Leads it is typically lower, 10 to 20%.
- **Decision Owner release levels.** A Decision Owner released at 60% contributes 60% to the pool, not 100%. Their remaining time belongs to their standing role and is not available.

The arithmetic is simple. Suppose the portfolio guild has eight potential Decision Owners. Two are on leave this week. One is already committed to an active Effort. Of the remaining five, each is released at 60%. The pool holds three available Decision Owner-equivalents at the committed release level, not eight. The same subtraction runs for Fleet Leads and for agent fleet capacity.

The output is a single statement, refreshed weekly: this many Decision Owners are available at the committed release level, this many Fleet Leads, this much fleet capacity, and therefore this many Efforts can be in flight. Three perspectives must be present when the Flow Council produces this statement: portfolio and product (what is most valuable), technology (what is feasible), and capacity management (what is genuinely available). An allocation made without any one of the three fails in a predictable way that takes weeks to surface and is expensive by the time it does.

Commitments are made against what exists, not against headcount. Headcount is a budget line. Capacity is what remains after leave, standing obligations, and partial release levels are subtracted.

The distinction does real work because headcount-based planning overpromises every single week. A portfolio that commits against headcount will start more Efforts than it can sustain, starve each one of attention, and produce the familiar symptom of everything being in progress and nothing being done. Planning against actual capacity prevents the overcommitment at the source.

Weekly

The pool is re-planned every week at Flow Council calibration. Capacity is a current fact, not an annual assumption.

Commit

Capacity goes to an Effort for as long as that Effort needs it. The commitment is to an Effort, not to a duration.

No team holds a standing claim on people or on fleet. When the Flow Council commits capacity, it names the Effort, the Decision Owner, and the people and fleet assigned. It does not name an end date, because some outcomes take four days and some take four months, and pretending otherwise produces either padded timelines or arbitrary deadlines that force scope cuts in the wrong places.

Duration-based commitment is the mechanism that makes every planning conversation a negotiation. "How long will it take" is an answerable question for routine work and a guess for everything else. When the commitment is to a duration, the guess becomes a contract, and the contract creates two bad incentives: pad the estimate to absorb uncertainty, or cut scope to hit an arbitrary date. Both waste capacity in ways that are invisible at the portfolio level.

Effort-based commitment removes the negotiation. The team runs daily cycles until the outcome is accepted or the Flow Council redirects capacity to something more valuable. The weekly calibration is where the Flow Council evaluates whether each active Effort still deserves its capacity, so the portfolio is never locked into a commitment it cannot revisit. Duration emerges from the work. It is not imposed on the work.

Commitment, in practice

A committed Effort has three things:

1. A named Decision Owner, released at the committed level, present daily. The No Owner No Now gate enforces this.
2. A team: at minimum a Fleet Lead and fleet capacity, sized to the work.
3. A place in the Now horizon, meaning it is funded, owned, and running daily cycles.

If any of the three is missing, the Effort is not committed. It sits in Next until all three are present.

Commitment is to an Effort, not to a duration. No standing claims. When the outcome lands, the capacity returns.

Release

The moment an outcome is accepted, the people and the fleet return to the pool, and the Flow Council redeploys them to the next Effort, usually within the same week.

Release is the half most organizations skip, and skipping it is what rebuilds standing teams.

Release sounds obvious and almost nobody does it. The reasons are always reasonable on their surface. The team has built context. The system might need follow-on work. The people are comfortable. The manager does not want to lose them. Each reason is true in isolation and collectively they guarantee that capacity never returns.

If capacity never returns to the pool, within two quarters you have permanent teams with new labels and the model has failed silently. The people are doing valuable work. The teams are functioning well. And the portfolio has lost the ability to move capacity to its highest-value outcome, which is the entire point of the model.

Make release explicit

Release is a minuted Flow Council action with a date. It is not something that happens by drift, by assumption, or by someone eventually noticing that the Effort ended three weeks ago.

The minute records four things:

1. The Effort that concluded.
2. The date capacity returns to the pool.
3. The people and fleet capacity being released.
4. Where they are redeployed, or that they return to the unallocated pool.

Making it a formal action with a record is what prevents the silent failure. If release requires a decision, someone has to make it. If it requires a record, someone has to notice when it does not happen.

Minuted

Release is a recorded Flow Council action with a date. Not a drift. Not an assumption. A decision.

How an Effort Ends

Release depends on someone declaring the outcome met, and Efforts drift when nobody owns that call. Two people are involved in ending an Effort, and that is not an accident.

The Decision Owner proposes closure

When the outcome's acceptance criteria are satisfied, the Decision Owner proposes closure. They are the person closest to the work, with the domain knowledge to judge whether the outcome has been reached. The proposal names the criteria that are met and points to the evidence in the harness and the ledger. It is a claim with proof, not an opinion.

The Flow Council confirms

The Flow Council confirms closure at weekly calibration. The council holds the portfolio view that the Decision Owner does not. An owner close to the work will sometimes keep going past the point of value, adding increments that improve what already works rather than releasing capacity to what has not started.

Two people are involved because the incentives diverge. The owner is invested in the outcome and naturally inclined to keep improving it. The council is invested in the portfolio and naturally inclined to redeploy capacity. Neither perspective alone produces the right call. Together they do. The confirmation is recorded with a date, and release follows immediately.

The one-week stall rule

An Effort that has produced no landed increment for a week is raised at calibration whether or not the owner proposes closure. The rule is automatic: the Flow Council

does not wait for someone to flag the pattern. It reads it from the evidence ledger at every weekly calibration and raises any Effort that has gone silent.

That pattern usually means one of three things:

- The outcome was reached and nobody declared it. This is the most common case. The team has been doing minor refinements that feel like progress but are not producing accepted increments, because the real work was done days ago.
- The Effort is blocked and the blocker was not escalated. The daily cycle stopped closing but the team absorbed the stall locally rather than raising it to the Flow Council.
- It should never have entered Now. The Effort was admitted with insufficient domain clarity, and the team has been specifying rather than building. Move it back to Next with a specific condition: "return when the outcome can be expressed as checkable criteria."

Each of these has a different resolution, but all of them start with the same action: the Flow Council names it, discusses it, and records a decision. Silence is the failure mode. A week without a landed increment is the signal that breaks the silence.

Fluid Formation

Fluid formation is a Flow Council decision rather than a team's, and it is narrower than the phrase suggests.

A team is fluid between Efforts and stable within one.

A team holds together for the whole life of an Effort. The Flow Council reforms it at the boundary. This distinction does real work because the phrase "fluid teams" is often heard as "people move constantly," which is the opposite of what the model intends.

Within an Effort, the team has accumulated spec history, domain context, and harness knowledge that does not survive a handover. Moving someone mid-Effort does not free capacity. It destroys context and forces the replacement to rebuild it, which costs more time than it saves.




Figure 2: Standing teams lock capacity to systems. The fungible pool lets capacity flow to outcomes. Arrows from Efforts back to the pool represent release, the step that keeps the model honest.

Three triggers for reformation

The Flow Council reforms a team on exactly three triggers:

1. **The outcome is accepted** and capacity releases back to the pool.
2. **A higher-value Effort outranks it at triage** and capacity is redirected. This is a portfolio decision, made deliberately and recorded.
3. **The Effort proves to need different expertise** than it appeared to need at formation. The team's composition no longer matches the work.

Non-triggers

Three things that feel like reasons to move people and are not:

- **Convenience.** Someone is available and another team could use help. If the first team's Effort is still in flight, the person stays.
- **Utilization smoothing.** The team had a slow day. Slow days happen inside healthy Efforts. They do not justify moving people.
- **A manager wanting someone back.** Standing obligations are accounted for in capacity planning. If someone's standing work was not subtracted at planning, fix the plan. Do not pull them from an active Effort.

Rebuilding context is expensive. A team mid-Effort holds accumulated spec history, domain context, and harness knowledge that does not survive a handover.

The Decision Owner is the strongest case for stability. If their release lapses, the Effort does not continue with a substitute. It returns to Next until sponsorship is restored, because the alternative is a proxy owner and a cycle that stops closing without anyone noticing.

Standing Teams Rebuild the Old Constraint

The argument for standing teams is familiar and sounds persuasive. Teams build context over time. Relationships improve with continuity. People perform better when they know their teammates. All true, and all studied. The research on team stability (Hackman, Huckman and Staats at HBS) consistently shows that stable teams outperform newly formed ones on the same task. The critical qualifier is "on the same task." The pool model does not dispute this. It asserts that the cost of context rebuild at Effort boundaries is lower than the cost of locking capacity to work that no longer needs it.

The decision framework is concrete. If the Effort will run for four or more weeks, team stability is worth the formation cost and the team holds together for the duration. If the Effort is done in under a week, the context rebuild cost is near zero because there is almost no context to rebuild. The pool model's value appears at the boundary between Efforts, not within them.

The problem is not with the argument for stability. It is with the consequence of permanence. If capacity never returns to the pool, within two quarters you have permanent teams with new labels. The portfolio can no longer move capacity to its highest-value outcome. And because the teams are stable and productive, nobody notices that the model has failed.

If capacity never returns to the pool, within two quarters you have permanent teams with new labels. The model has failed silently.

This failure is silent because it produces no crisis. The teams are doing good work. The people are happy. The systems are maintained. But the portfolio has lost its ability to respond to a change in priority, and the next time a high-value outcome appears, the organization will discover that all of its capacity is locked to things that matter less and none of it can move.

That is the old constraint, rebuilt under new vocabulary.

The mechanism of silent failure

It happens in five stages, each one reasonable, each one harder to reverse than the last.

Stage 1. An Effort ends but the team stays together because follow-on work might appear. Nobody records a release decision. The pool does not notice a missing return because no return was expected on a specific date.

Stage 2. Follow-on work appears, because work always appears for a team that is sitting together and looking for something to do. The work is real. It is also work that the portfolio did not prioritize through triage. It entered the active portfolio through the side door of team availability rather than through the front door of the No Owner No Now gate.

Stage 3. The team begins to identify with a system rather than with an outcome. They become the people who know that system, and the organization begins to route work to them based on the system it touches rather than on portfolio priority. Knowledge of the system becomes the justification for keeping the team together, which is circular: the team has the knowledge because the team was never released.

Stage 4. The Flow Council stops planning capacity for that team because it is understood to be spoken for. The pool shrinks by the size of that team, and the remaining Efforts compete for whatever is left. Nobody decided to exempt the team from the pool. It happened by drift.

Stage 5. A new, high-value initiative arrives that needs capacity the portfolio does not have, because the capacity is locked in teams that cannot be disturbed. The organization hires, which takes months, and the initiative waits. The portfolio now carries both the cost of the standing teams and the cost of the delay, and nobody connects the two.

The correction

The correction is not complicated. It is uncomfortable.

Make release a recorded, minuted, non-negotiable Flow Council action. Every Effort that ends produces a release decision with a date. Every person and every fleet allocation returns to the pool. The Flow Council redeploys them, usually within the same week, to whatever the portfolio needs most. In the portfolios where we have enforced this, the median time from outcome closure to redeployment is three working days. Without enforcement, the median is “never,” because the team stays together indefinitely and the redeployment conversation never happens.

The discomfort is real. People will feel unsettled by not belonging to a permanent team. Managers will feel the loss of control over “their” people. The Flow Council will face pressure to let teams stay together. Hold the line anyway, because the alternative is the silent reassembly of the structure the model was built to replace.

Failure Modes

Two failure modes directly threaten the pool model.

Capacity never released

Tell: The same people remain on the same Effort beyond the point where the outcome was reached.

Correction: Make release a minuted Flow Council action with a date. If the one-week stall rule is enforced at calibration, this failure becomes visible before it compounds. If the rule is not enforced, the silent failure described above runs to completion within two quarters.

Trains reassemble

Tell: Flow Council members advocate for “their” teams rather than for the portfolio. Capacity decisions start reflecting which council member speaks loudest rather than which outcome the portfolio needs most.

Correction: Change how Flow Council members are evaluated. If they are evaluated on the outcomes their area produces, they will optimize for their area. If they are evaluated on the portfolio’s throughput and outcome delivery, they will trade their area’s interest for the portfolio’s. The incentive change must be explicit and visible, because the old incentive reasserts itself quietly.

How the Pool Model Holds Together

- **The No Owner No Now gate.** This controls what enters the pool. Without it, capacity commits without ownership and the pool fills with unfunded good ideas. The gate and the release are two halves of the same discipline.
- **Decision Owner qualification.** These screens determine whether a named person is fit to own an Effort. The pool model depends on real owners, not proxies, because a proxy prevents the daily cycle from closing and the Effort from reaching the point where capacity can be released.

Is your capacity truly fungible?

If you cannot name the last time capacity was released back to the pool, the model has already failed. 30 minutes to talk about what to do about it.

[Book a call](#) | singularics.ai

SINGULARICS

We help large organizations close the intent gap between strategy and execution so that when they accelerate with AI, they accelerate in the right direction.

singularics.ai | [Book a conversation](#)