



White Paper

The One-Day Cycle

Hour by hour. How the daily loop of Specify, Build, Judge, Land turns once per day for as long as the outcome takes.

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 Loop, Not a Sequence	5
1.1 Empowerment at each step	6
2 The Day, Hour by Hour	7
2.1 Before the day starts	7
2.2 First hour: decomposition	7
2.3 Middle of the day: build	7
2.4 Late afternoon: prep	8
2.5 End of day: close	8
3 The Day-Ahead Tension	9
3.1 Most of tomorrow's spec is already known	9
3.2 A shallow queue, not a fixed lead time	9
4 Daily Planning	10
4.1 The five-item agenda	10
4.2 Outputs	11
4.3 The tell that it has gone wrong	11
5 Confirm, Amend, Replace	12
5.1 Confirm (60–70% of days)	12
5.2 Amend (20–30% of days)	12
5.3 Replace (5–15% of days)	13
5.4 The Replace rate as a diagnostic	13
6 When a Day Fails	14
7 A Red Harness Is Not an Acceptance Decision	15
7.1 When the criterion is wrong	15
8 The Working Week	16
9 Failure Modes	17
9.1 Slices too big	17
9.2 Demo-only acceptance	17
9.3 Spec as afterthought	17
10 Mechanisms That Make the Cycle Work	18


Abstract

The AI-native operating model runs on a single repeating unit: the one-day cycle. Specify, Build, Judge, Land. It turns once per day for as long as the outcome takes, whether that is four days or four months. Judge sits inside the loop rather than near the end, because judgment is not an approval gate applied after the work is finished. It happens every day, on every increment, by the person who has been present the whole time. Daily planning closes each turn and opens the next. This paper walks through the day hour by hour, explains how the spec runs a day ahead without requiring certainty, defines the three calls available at daily planning, and establishes what a failed day looks like when the model is working correctly.

The Loop, Not a Sequence

The single most common reaction to this model from experienced delivery leaders is: “A one-day cycle is waterfall compressed into 24 hours.” It is not. The difference is structural. In a waterfall, the spec authors hand off to the builders, who hand off to the testers, who hand off to the approvers. In the one-day cycle, the same Decision Owner who writes the spec judges the result, and they are present through the build. There are no handoffs. There is one loop with one owner. That distinction is what separates a daily cycle from a miniature phase gate.

The work lifecycle has a daily zone that repeats: Specify, Build, Judge, Land. It is a loop and not a sequence of gates. It turns once per day for as long as the outcome takes, which may be four days or four months.



SINGULARICS Copyright © 2026 | singularics.ai

Figure 1: The daily loop. Specify, Build, Judge, Land, then Daily Planning closes the turn and sets the next spec. The cycle turns once per day for as long as the outcome takes.

Reading these four steps as a sequence introduces two misreadings the model cannot survive.

The first is that Judge appears as a late step, which makes it read as an approval gate. Judgment is not applied to the work after it is finished. It happens every day, on every increment, by the person who has been present the whole time. If Judge sits outside the loop, it becomes a queue, and the queue is what this operating model exists to remove.

The second is that Specify reads as a phase that precedes delivery. It is not. It is the first move of every single day. The spec is not a plan for the project. It is a contract for the day.

Daily planning closes each turn and opens the next. It is the hinge that makes the four steps into a loop rather than a line. Without it, each day is an isolated event rather than a turn of a cycle that accumulates learning.

If the Decision Owner cannot accept, they are not an owner, and the daily cycle does not close. The cycle is the cycle, and there is no version of it that closes weekly.

Empowerment at each step

Step	Who decides
Specify	Decision Owner, with a spec agent and a domain expert
Build	Fleet Lead
Judge	Decision Owner
Land	Decision Owner, inside the clearance rules

Read down that column and the empowerment structure is visible. The Decision Owner owns the entire daily loop, which is the concrete meaning of delegation. The Fleet Lead owns how the work is built rather than whether it is right.

The Day, Hour by Hour

A realistic schedule with time budgets. The times are indicative, not prescribed, but the proportions hold. The spec makes the day's contract explicit, and the harness turns its acceptance criteria into evidence on every push. Those two instruments make this schedule possible.

Block	Time budget	Who is present
Spec ready (day before)	45–90 min	Decision Owner, spec agent, domain expert
Decomposition	30–60 min	Fleet Lead, fleet, Decision Owner on call
Build	4–5 hours	Fleet Lead, fleet
Decision Owner available	2–3 hours	Decision Owner (responsive, not scheduled)
Close: demo + acceptance	15–20 min	Decision Owner, Fleet Lead, stakeholders
Close: daily planning	20 min	Decision Owner, Fleet Lead, Flow Council lead

Before the day starts

The spec exists, reviewed and signed. The domain context the increment needs has been injected by the Domain Knowledge Network. Nothing about the day's direction is open. The build day begins ready because the spec was written the day before. If the spec is not signed, the build day does not start. This is a hard gate, not a courtesy.

The Decision Owner spent 45 to 90 minutes on the spec the previous day: correcting the draft the spec agent produced, adding the constraints only a domain person would know, and signing every acceptance criterion. If they spent three hours, the slice is too big.

First hour: decomposition

The Fleet Lead reads the spec with the fleet, decomposes it into work the agents can execute, and confirms the harness criteria are expressible as checks. If a criterion cannot be checked, that is raised now and not at four in the afternoon. The Decision Owner is reachable for this window specifically, because questions about intent surface here more than anywhere else in the day. Budget 30 to 60 minutes.

Middle of the day: build

The fleet implements against the contract. The harness grows alongside, since agents write and maintain coverage as they build. The Fleet Lead directs, reviews at altitude, and course-corrects. The build block is 4 to 5 hours, which is the productive core of the day.

The Decision Owner answers questions as they arrive during this block, typically three to six over the day, each taking minutes rather than a meeting. Available means reachable within minutes rather than schedulable within days. They do their standing business work during this block but treat team questions as interrupts with priority.

4 hrs

of burned build time is the cost when a Fleet Lead waits a full afternoon for a single judgment call. This happens when the Decision Owner's "available" block is cut for other meetings. Protect that block above all others.

The decision rule for the Decision Owner during the build block is simple. If the team's question can be answered in under five minutes, answer it now. If it requires more than five minutes of thought, give the Fleet Lead a provisional answer with the constraint that must hold, then provide the full answer within two hours. If the question falls outside the Decision Owner's envelope, escalate to the Flow Council immediately rather than letting the team wait.

Late afternoon: prep

The harness runs clean, or it does not. The team prepares the demo against the criteria in the spec, not against a happy-path walkthrough. The demo shows stakeholders the work in a form they can react to. What the demo cannot do is carry the acceptance decision, because a demo shows one path through the system while the harness checks every path, every time.

End of day: close

Demo, judgment, acceptance or a named blocker, landing, then daily planning. Budget 45 to 60 minutes for all of it combined. This is the most important block in the day because it is where the cycle closes and the next one opens.

The sequence within this block is deliberate. The demo comes first, 15 to 20 minutes, so everyone sees the same evidence. The judgment and acceptance decision follow immediately, made by the Decision Owner against the harness results and the spec criteria. If the increment is accepted, it lands. If it is not, the blocking criteria are named precisely. Then daily planning runs while the evidence is fresh, 20 minutes.

45–60 min

for the close: demo, judgment, acceptance, landing, and daily planning combined. This block is not optional and not compressible.

The Day-Ahead Tension

State the tension plainly. Spec runs a day ahead so the build day starts ready. But the Decision Owner often does not know what they want next until they have seen today's demo. Those two facts appear to contradict each other.

They do not.

Most of tomorrow's spec is already known

Separate every spec into two parts.

The stable part is the outcome, the constraints, the policy envelope, the data model, the acceptance philosophy, and the general shape of the next several increments. This is written days or weeks ahead and changes rarely.

The volatile part is which specific slice comes next and its exact criteria. This is what the demo informs.

70–80%

of the spec is stable context, written days or weeks ahead. The remaining 20–30% is the volatile slice decision that daily planning resolves.

Once you see it this way, daily planning is not writing a spec from scratch in twenty minutes. It is making a small delta decision against a mostly-written contract. The 70–80% stable portion is the reason a Decision Owner can produce a signed spec in 45 to 90 minutes rather than three hours: most of the document already exists and was not invalidated by today's demo.

A shallow queue, not a fixed lead time

Do not think of it as "the spec is written exactly one day early." Think of it as a shallow queue of one to three candidate specs, kept ready.

Most days the demo confirms the direction and you pull the next queued spec unchanged. Some days the demo surprises you and you reorder the queue or amend the spec at the top. The queue absorbs volatility that a rigid one-day lead time cannot.

Keep it shallow. A queue deeper than three specs is a backlog forming, and it means you are speculating further ahead than your learning rate justifies.

If every demo overturns the plan, your slices are too big. A demo that invalidates a week of planned work means you planned a week ahead of your knowledge. A demo that adjusts tomorrow means you sliced correctly.

Daily Planning

Timebox. 20 minutes. It is a decision forum, not a status meeting.

Attendees. Decision Owner, Fleet Lead, and a Flow Council lead. The domain expert attends when the next increment touches their specialty, which is perhaps two days in five.

Position in the day. Immediately after the demo and the acceptance decision, while the evidence is fresh.

The five-item agenda

Twenty minutes, five items, strict time allocations. The allocations add to 18 minutes, leaving two minutes of margin. If an item threatens to run long, the Flow Council lead calls it and moves on.

1. **What landed today, and where it landed.** One minute. A factual statement, not a discussion. "Increment 14 accepted, landed in staging" or "not accepted, criteria 3 and 7 fail."
2. **What the demo and the harness told us that we did not know this morning.** Five minutes. This is the substance of the meeting. If the team learned nothing, either the slice was trivially predictable or nobody was paying attention.
3. **Confirm, Amend, or Replace.** Five minutes. The single decision about tomorrow. One word decides the call; the remaining time covers what the Amend changes or what the Replace requires.

4. **Which domain expert tomorrow's spec needs, and whether they are available.** Two minutes. If the expert is not available, the spec is adjusted to avoid the domain it would have touched, or the Effort pauses for a day. It does not proceed without the domain input it needs.
5. **Anything blocked that the Flow Council must clear tonight.** Five minutes. Blockers are stated as specific questions with a named owner. "We need a decision on X from the Flow Council by tomorrow morning" rather than "we are kind of stuck."

20 min Five items, strict allocations. If it runs past 30 minutes, it has become a working session. Split it.

Outputs

Every day, without exception, daily planning produces three things:

- A named decision on tomorrow's spec: Confirm, Amend, or Replace.
- A named domain expert if one is needed, confirmed as available.
- A named blocker with an owner and a deadline, or an explicit statement that there are none.


If daily planning is not producing these outputs, it is not daily planning. It is a conversation.

The tell that it has gone wrong

If daily planning routinely runs past 30 minutes, it has become a working session. Split it. Decisions happen in the 20 minutes, and spec authoring happens afterward with only the people who write it.

Confirm, Amend, Replace

At daily planning, after the demo, the group makes exactly one decision about tomorrow. There are three calls, and only three.



SINGULARICS Copyright © 2026 | singularics.ai

Replace rate is a diagnostic. Persistently above 30% means slices are too big.
Persistently at 0% means the work may not need this model.

Figure 2: Expected distribution of the three calls. The Replace rate is a diagnostic: persistently above 30% means slices are too big; persistently at 0% means the work may not need this model.

Confirm (60–70% of days)

The demo went as expected. Pull the next queued spec as written. No changes, no adjustments. The direction was right and the next increment follows from it naturally. This should be the majority of days, because a well-sliced outcome produces increments that build on each other predictably.

A Confirm day takes no additional spec work. The spec is already signed. Tomorrow's build starts on time. This is the reward for maintaining a well-curated queue.

Amend (20–30% of days)

The demo revealed something that changes the next increment's criteria but not its direction. The Decision Owner adjusts the top spec that evening or first thing the next morning, typically 20 to 30 minutes of work. This is the most common learning response: the direction holds but a detail shifts.

Typical Amend triggers: a data shape was different than expected, a policy constraint was stricter than assumed, or the demo revealed a user need that adjusts the next slice's acceptance criteria without changing what the slice is about.

Replace (5–15% of days)

The demo invalidated the plan. The queued spec is wrong and a new one is needed. The direction itself changed, not just the criteria. This is the strongest learning signal the cycle produces.

Replace days are not failures. They mean you learned something expensive early, which is the entire point of shipping daily.

When a Replace happens and the new direction is not yet clear, do not force a build day. Run tomorrow as a spec-and-explore cycle. The fleet investigates, prototypes throwaway options, and gathers what the owner needs to decide. The day still ends with something in someone's hands, but the artifact is a decision rather than an increment. This is the "learn" branch of ship-or-learn, and it is a legitimate, planned outcome rather than a lost day.

The Replace rate as a diagnostic

Track the Confirm/Amend/Replace ratio weekly. It is the earliest signal of slicing health and spec quality.

Replace persistently above 30%. Your slices are too big. Each demo is over-turning a plan that assumed more than one day's worth of knowledge. The correction is to halve the slice and re-run the checkability test on the criteria. If the work is genuinely exploratory, acknowledge it: change the Effort's cadence to spec-and-explore cycles until the direction stabilizes, then return to normal build days.

Replace persistently at 0%. You are building something so well understood that you should question whether it needed this model at all. Predictable work is better served by simpler machinery.

Amend above 40%. The stable portion of the spec is not stable enough. The Decision Owner is rewriting too much context between days, which means the outcome brief or the policy envelope was not established clearly enough at the start.

A healthy ratio is a sign that the team is learning at the right rate: fast enough to catch mistakes early, slow enough that the direction mostly holds.

When a Day Fails

A failed day is one where the increment is not accepted. This is normal and should happen regularly. The requirement is that the failure is precise.

*The unacceptable end to a day is "it's not quite there yet."
If a team ends days vaguely, the spec was vague. Go back
to the checkability test.*

There are exactly three acceptable ways for a day to end:

1. **Accepted and landed.** The harness is green. The Decision Owner judges the increment against the criteria, accepts it, and it reaches its destination: production if the clearance path allows it, a hands-on environment if not. The acceptance is recorded in the evidence ledger with the harness results attached.
2. **Not accepted, with specific failing criteria named.** The harness shows which criteria fail and here is the evidence. "Criterion three fails because the error response does not include the field name. Criterion seven fails because the calculation uses the wrong rate." That precision is what makes tomorrow's fix tractable. The failing criteria become the first items in tomorrow's spec.
3. **Decision required outside the Decision Owner's envelope.** The increment surfaced something the Decision Owner cannot resolve inside their bounded authority. The exact question for the Flow Council is stated as a specific, answerable question, not a status update. "Should we apply the 2024 or 2025 rate table for this counterparty class?" rather than "we are not sure about the rates." The blocker is raised that evening and resolved within one cadence.

No other ending is acceptable. "It's not quite there yet" is not a judgment. It is a deferral. "Most of the criteria pass" is not an acceptance. It is a partial, and partials do not land.

When a team ends days vaguely, the spec was vague. The correction is not to try harder at acceptance but to go back to the checkability test on the criteria. If the criteria do not support a clean yes or no, the spec failed before the build started.

A failed day that names its failures precisely is more valuable than a successful day that lands something nobody inspected closely. The first produces a correction. The second produces drift.

20–30%

of days ending without acceptance is normal in the first month. If every day succeeds, the slices are trivially small or the acceptance bar is too low. If more than 40% fail, the specs are too ambitious or the Decision Owner is not present enough. Both extremes are diagnostics. Aim for 70–80% acceptance over any four-week window.

A Red Harness Is Not an Acceptance Decision

The most common way the daily cycle degrades is a team accepting an increment on a red harness because the failures look cosmetic. If the harness is red at close, the increment does not land. The day ends with the failing criteria named, and that is the blocker branch working as designed.

This rule is simple and it must be held without exception, because the exceptions compound.

An increment waved through on a red harness teaches the team that the harness is advisory rather than authoritative. Within a month, red becomes normal. Within two months, the harness stops being maintained. Within three, the Decision Owner is accepting on demo alone, and the entire evidence layer the model depends on has eroded.

When the criterion is wrong

There is a legitimate case where a criterion is failing and the build is right. The criterion itself was wrong. It specified the wrong behavior, the wrong threshold, or the wrong constraint. The build did the right thing and the check is testing for the wrong thing.

This is a spec correction. It is made in daily planning, not waved through at acceptance. The Decision Owner acknowledges that the criterion was wrong, the spec is amended, the harness is updated, and the increment is re-evaluated against the corrected criteria.

A criterion failing because the criterion is wrong is a spec correction in daily planning, not a waiver at acceptance. The distinction does real work because the second habit is how a harness stops meaning anything within a month.

The distinction does real work because the two habits look similar in the moment and diverge completely over time. Correcting the spec maintains the integrity of the harness. Waving through a red result destroys it.

The Working Week

The cycle runs on working days. A four-day week from a holiday is four cycles, not a compressed five. A three-day week is three cycles. The spec queue absorbs the gap without anyone re-planning.

Teams that try to make up a lost cycle by doubling a spec produce the oversized increments that break the model. A two-day spec crammed into one day produces an increment too large to judge cleanly, criteria too numerous to check in a single harness run, and a demo too sprawling to evaluate in thirty minutes. The result is either a vague acceptance or a vague rejection, and both are worse than simply running one fewer cycle that week.

The temptation to compress is strongest before a deadline. Resist it. The deadline does not move because you crammed two days of work into one. It moves because you shipped something every day and the scope was managed at daily planning. If the scope was not managed at daily planning, compressing the last week will not save it.

A four-day week from a holiday is four cycles, not a compressed five. No make-up days. The spec queue absorbs the gap.

The working week also has no special ceremony attached to its boundaries. Monday is a cycle. Friday is a cycle. There is no sprint start, no sprint end, no retrospective that spans the week. The unit of rhythm is the day, and the week is simply how many days it contains.

If the team needs a retrospective or a longer reflection, schedule it as its own event. Do not load it onto the daily planning of the last day of the week, because that stretches the 20 minutes into an hour and teaches the team that Fridays are different. They are not.

The practical difference from a sprint is worth stating plainly. A sprint is a time container with a planning ceremony at the start and a review at the end. The one-day cycle has no container. It has a loop. Planning happens every day. Review happens every day. There is no batch, no ceremony accumulation, and

no two-week delay between a question and its answer. Organizations that have operated both report that the daily cycle surfaces problems in hours that the sprint cadence surfaced in weeks.

Failure Modes

Three failure modes threaten the daily cycle directly.

Slices too big

Tell: The Replace rate is persistently above 30%, or increments regularly span more than one day.

Correction: Halve the slice. Re-run the checkability test on the criteria. If the spec has more than about fifteen acceptance criteria, it is two specs. If the demo takes more than 20 minutes to walk through, the increment was too large to judge cleanly in a single session.

A two-day increment is not a minor overshoot. It means the spec was underspecified, the fleet was under-directed, or the outcome is more complex than the team acknowledged. Each of those is a different problem, but the symptom is the same: the cycle did not close.

Demo-only acceptance

Tell: Acceptance happens because the demo looked right, without reference to the harness results or the spec criteria.

Correction: Return to evidence. A demo shows one path through the system. The harness checks every path, every time. The rule is simple: if the harness is red, the increment does not land, regardless of how good the demo looked. If acceptance is routinely made on the demo alone, the harness will stop being maintained within a month, and the evidence layer the model depends on will erode within three.

Spec as afterthought

Tell: The team starts building while the spec is still open, or the spec is signed after the build day begins.

Correction: Hard gate. No build day starts without a signed spec. The spec is the contract the fleet builds against and the instrument the Decision Owner accepts against. Without it, the team is building toward an undeclared target, and

the end-of-day judgment is a subjective reaction rather than a decision against evidence.

Mechanisms That Make the Cycle Work

- **The spec.** This is the contract the daily cycle builds against. Its authorship split, checkability test, and one-day scope make a daily acceptance decision possible.
- **The harness.** This is the evidence layer the acceptance decision is made against. It runs on every push in CI, not before demos, and it is what allows a business domain expert to accept work every day without reading code.
- **The No Owner No Now gate.** This ensures the daily cycle has a Decision Owner present to close it. Without the gate, cycles run without acceptance and increments accumulate without judgment.

Ready to run the daily cycle?

30 minutes to walk through the mechanics. No pitch.

[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](#)