10

WHITE PAPER · PORTFOLIO & DELIVERY GOVERNANCE

# Decision Friction Is a Portfolio Risk

Make consequential choices visible before delay becomes rework

**Marco Policani**

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

July 2026

EXECUTIVE SUMMARY

**Decision friction is managed when the owner, evidence, and latest useful date are visible.**

Portfolios govern work in motion while the decisions controlling that work live in inboxes, meetings, and memory — where delay is invisible until it becomes rework, idle capacity, or an accidental choice. This paper builds the decision inventory: consequential choices named, owned, evidenced, dated, connected across investments, and closed.

**01** Give every consequential decision a latest useful date — the point where the answer stops mattering — and escalate against it as ordinary practice.

**02** Test every analysis request against the alternative it could eliminate; the rest is scheduling cover.

**03** Convert recurring decisions into standing policy, the healthiest way a decision register shrinks itself.


## The argument

Portfolios track work in motion but often leave the decisions controlling that work scattered across meetings, inboxes, and individual memory. A decision inventory makes delay governable before ambiguity becomes rework, idle capacity, or an accidental choice.

The schedule slip that taught me this pattern looked, from the portfolio view, like a delivery failure. A program was eight weeks late on an integration milestone, and the retrospective was preparing the usual vocabulary: estimation, complexity, vendor performance. Then the delivery lead walked me through the actual eight weeks. The team had spent five of them waiting to learn which of two data-retention policies applied to the records they were migrating. They had asked in week one. The question had traveled through a compliance mailbox, two steering committees, and an executive's chief of staff, gathering interpretations at each stop. No one had refused to decide. The decision simply had no address.

Once you learn to see this, you see it everywhere. A milestone can be late because a team executed poorly. It can also be late because the organization never settled a choice the team could not make for itself: which policy applies, what tradeoff is acceptable, who owns the process being changed, whether a dependency will be funded, which of two outcomes takes precedence when they conflict. Conventional reporting discovers the difference only after schedule variance appears, and usually attributes it wrong.

Decision friction is the elapsed time, repeated translation, and organizational energy consumed between recognizing that a choice is needed and making it in a form the work can use. This paper, written for portfolio executives, PMO leaders, sponsors, and program leaders, treats that friction as a first-class portfolio risk and proposes the instrument for governing it: a decision inventory with six disciplines. PMI's complex-project research and practitioner IT-governance writing both connect delivery outcomes to decision-making quality [1][2]; the inventory design comes from the portfolios where I learned to count the waiting.

A note on what this paper is not arguing. Faster is not always better; the objective is decisions that arrive inside their window of usefulness, which sometimes means slowing a choice down to let real evidence land. Friction, as used here, is not deliberation. It is the untracked interval in which nobody is deliberating and nobody is deciding — the question is simply in transit, and the meter is running.

## Why decisions escape governance

Work is easy to govern because work is visible: it has a name, a plan, a budget line, and a person who will be asked about it on Thursday. A pending decision has none of those properties by default. It exists as a question in flight — mentioned in a status note, parked in a risk log, alive mainly in the worry of the people waiting on it. The portfolio governs artifacts it can enumerate, and nobody enumerates the questions.

Escalation culture makes this worse in a particular way. In most organizations, escalating a decision reads as either failure or aggression — the team could not handle it, or the team is going over someone's head. So teams delay escalation past the point where it would have been cheap, try workarounds first, and phrase the eventual escalation so diplomatically that it arrives without a deadline or a requested answer. The question enters the executive layer as context rather than as a decision, and context does not get decided.

The cost structure is what makes friction a portfolio matter rather than a nuisance. Waiting teams do not simply pause; they burn capacity on translation — re-explaining the question to each new forum — and on speculative work built on guessed answers, some of which becomes rework when the real answer lands. Worst, options expire silently: while the choice waits, procurement windows close, architects get reassigned, and the eventual decision is made by default among whatever alternatives survived. An accidental choice, made by calendar.

## The portfolio decision inventory

The inventory applies to consequential choices the discipline that schedules apply to work: each one named, owned, evidenced, dated, connected, and closed. Six disciplines, described in sequence — with the failure each prevents and the evidence that shows it operating.

EXHIBIT 1

### Six disciplines make consequential choices governable

1. **State the decision**
   A decision record begins with one answerable choice, not a topic for discussion.
2. **Name the accountable owner**
   The owner must have authority over the choice or a defined route to the person who does.
3. **Specify the evidence need**
   The inventory should show what is missing and whether it can realistically change the choice.
4. **Set the latest useful date**
   A due date describes administration; a latest useful date describes when the decision can still shape the work.

5 **Expose downstream risk**

Decision delay should be connected to capacity, dependencies, cost of change, and benefits—not hidden inside schedule commentary.

6 **Close the loop**

A decision is complete only when its rationale, conditions, and implications reach the people and records that govern execution.

## State the decision

An inventory entry begins with one answerable choice, not a topic. 'Architecture concerns' is a theme; 'choose between extending the current billing platform or adopting the vendor module for usage-based pricing, by March, for the mid-market segment' is a decision. The difference decides whether an executive forum can act in one sitting or must first spend a meeting discovering what is being asked.

Vague framing is not laziness; it is often strategy. A broad theme lets every party hear their preferred question, delays the moment when someone must lose, and protects the requester from proposing an answer that might be rejected. The inventory's first discipline breaks that equilibrium: entries must carry the alternatives actually available, a recommendation with its reasoning, and the consequence of each path in operating terms.

The gatekeeping works both directions. Items that arrive as themes get returned for framing before they consume forum time — a service the PMO can perform in an hour that saves the forum a month. And well-framed items earn fast handling, which teaches requesters that framing pays. Within two quarters, the quality of what reaches the forum changes, because the market for its attention finally has a price.

## Name the accountable owner

Every entry names the person who can bind the organization to an answer — or the defined route to whoever can. Ownership here means authority over the contested boundary, not proximity to the topic. The distinction shows up the moment a real conflict arrives: a decision 'owned' by a director whose authority stops at her function will be made, contested, and remade when the neighboring function objects.

Ownership discipline changes requester behavior too. Teams that know an entry must carry a bindable owner stop lobbing questions upward as open letters and start doing the homework of locating authority — which frequently reveals that the authority is theirs. A meaningful fraction of 'escalations' in any organization are actually requests for permission the requester already holds, made for political cover. The inventory's ownership field quietly returns those to their rightful place.

The common anti-patterns are the committee-as-owner and the distribution-list-as-owner. A committee can hold decision rights if its charter says so and quorum rules exist; a committee that merely discusses is a venue, not an owner. A distribution list is nobody. The inventory forces the question the org chart avoids: for this specific boundary, who decides? Asking it once, in calm conditions, is far cheaper than discovering the answer during a dispute.

Where genuine authority gaps appear — choices that fall between charters, or reserved matters nobody reserved — the entry routes to the escalation path, and the gap itself gets recorded. A quarter's accumulation of authority gaps is a governance redesign agenda, delivered as evidence instead of opinion.

## Specify the evidence need

The third discipline separates decisions waiting on evidence from decisions hiding behind it. Each entry lists the decisive facts and uncertainties: what is known, what analysis is underway, and — the critical field — whether any realistic new information would actually change the choice. Some questions are hard because data is missing. Others are hard because every answer costs someone something, and more analysis is simply the socially acceptable form of not deciding.

The request for another study is the classic deferral instrument, and it works because it sounds like rigor. The inventory counters with a simple test attached to each evidence request: which alternative would this analysis eliminate, and would its cost and delay be justified by that elimination? Analysis that cannot name the alternative it might kill is not evidence work; it is scheduling cover.

Where uncertainty is real, the entry records it honestly — ranges, dissent, the residual unknowns — and the decision proceeds on bounded uncertainty when the cost of further waiting exceeds the value of further knowing. That sentence is easy to write and culturally hard to live; the inventory's contribution is making the tradeoff explicit each time, so deciding under uncertainty becomes a recorded practice rather than an act of individual courage.

## Set the latest useful date

A due date says when an answer is requested. A latest useful date says when the answer stops mattering — the point past which downstream commitments will have foreclosed the alternatives and the decision will have made itself. The two dates are different instruments, and the second is the one that governs friction.

Deriving it is a planning exercise, not a guess: walk backward from the first downstream commitment that depends on the answer. When must procurement start to hold the vendor window? When does the architect roll off? When does the parallel build reach the fork where the two paths diverge in cost? The latest useful date is the earliest of those, minus the time the decision itself needs. Teams can almost always produce it in an afternoon once someone asks.

The date transforms escalation politics. 'We need this decided' is a complaint; 'this choice loses its cheapest alternative on April 12' is a fact with a clock. Escalation triggered by an approaching latest useful date is not aggression — it is the system working, and forums learn to read it that way when the dates prove reliable. The measure that matters: decisions made inside their latest useful dates, and the trend of that rate.

## Expose downstream risk

The fifth discipline aggregates. A single delayed decision is a project problem; the same policy question blocking three programs is a portfolio problem wearing three disguises. Because each affected team reports its own local symptom — one waits on 'compliance input,' another on 'data guidance,' a third logs a schedule risk — the shared bottleneck stays invisible unless some view joins them.

The inventory joins them by linking each decision to the investments, constrained roles, and dependencies it touches. The aggregation views practically write the portfolio review agenda: decisions blocking the most value, owners who appear most often, subjects that recur quarter after quarter, queues forming in front of particular forums. Each is a systemic finding no individual project could have surfaced.

The aggregation also reprices escalation for executives. A leader who sees one delayed decision sees a team's impatience; a leader who sees that decisions in one subject area have averaged forty days against a fourteen-day useful window for three quarters sees a broken mechanism they own. The inventory does for decision latency what incident data did for reliability — it converts anecdotes into a managed distribution.

Recurrence deserves special attention, because a repeating decision is an unwritten policy asking to be written. The same data-residency question, decided case by case five times, is a standing policy waiting to be written; the same vendor-versus-build fork recurring across programs is an architecture principle nobody has stated. Converting recurring decisions into standing rules is how the inventory shrinks itself — the healthiest trend a decision register can show.

## Close the loop

A decision is complete when the people executing under it are working from the answer — not when the meeting ends. The gap between those two moments is its own species of friction: leaders decide, minutes go out, and three weeks later a team is still building against the old assumption because the decision reached a distribution list rather than the plan. Different functions can even hold different understandings of what was decided, each traceable to the same ambiguous minute.

The distribution problem is mechanical and solvable; the interpretation problem is subtler. Two functions can leave the same meeting with compatible-sounding, materially different readings of the decision — 'we approved the vendor module' versus 'we approved the vendor module pending the security review' — and each will build on its version until the versions collide. Closure discipline exists for exactly this: the recorded answer, in operational language, is the version, and everything else is memory.

Closure is a checklist, and it is short: the answer recorded with its rationale and conditions; affected owners acknowledging receipt; plans, budgets, and boundaries updated where the answer changed them; and the reopening condition stated — what changed circumstance would legitimately put this question back on the table. That last field prevents both bad futures: decisions relitigated whenever a loser finds a new forum, and decisions treated as eternal after their premises died.

The rationale field earns its cost the first time context shifts. A recorded 'we chose the vendor module because volumes were under the threshold' converts a later volume surge from an awkward reversal into a clean, pre-agreed reopening. Decisions with reasoning attached age into institutional knowledge; decisions without it age into folklore.

## One entry, end to end

A complete entry is compact enough to show whole. Decision: adopt the vendor's usage-based billing module for mid-market, or extend the current platform. Owner: VP Revenue Technology; escalation to the investment committee if the answer changes this year's platform budget. Evidence: vendor reference checks complete;

volume projections in hand; open uncertainty is migration effort, sized by a two-week spike now underway — the only analysis that could still change the answer. Latest useful date: April 12, when the vendor's contract-year pricing lapses and the parallel design work forks. Downstream: blocks the pricing pilot, holds two integration engineers idle from May, touches the CPQ replacement. Recommendation: adopt, contingent on the spike confirming migration under 400 hours.

That is eleven lines. It took the requesting team perhaps two hours to assemble, most of it the backward walk to April 12. A forum receiving that entry can decide in fifteen minutes or delegate it on the spot — and either outcome is governance working. The same question, arriving as 'we need alignment on billing architecture,' would have consumed three meetings just becoming askable.

The entry also shows what the inventory does not demand: no lengthy business case, no scoring model, no new template ecosystem. The disciplines are a quality bar for a conversation the organization was already trying to have, badly.

## The disciplines as one instrument

The six disciplines are one workflow read in order: frame it, own it, evidence it, date it, connect it, close it. Weakness travels backward through the chain — closure fails when ownership was vague, dates mean nothing when framing was thematic, aggregation lies when entries are incomplete. When the inventory disappoints, diagnose from the front: nearly every failure is a framing or ownership failure that surfaced downstream.

EXHIBIT 2

### What each discipline resolves

**State the decision**
Return vague items for framing before they consume executive review time.

**Name the accountable owner**
Escalate promptly when the named owner cannot bind the organization.

**Specify the evidence need**
Decide with bounded uncertainty when further analysis costs exceed its decision value.

**Set the latest useful date**
Escalate before optionality disappears rather than after the milestone turns red.

**Expose downstream risk**
Resolve systemic bottlenecks at portfolio level instead of asking each team to chase the same answer.

**Close the loop**
Reopen the decision only when a named condition changes, not whenever disagreement resurfaces.

The inventory coexists comfortably with decision-rights frameworks — RACI charts, delegation matrices, and their relatives — and fixes their chronic weakness. Those frameworks describe who should decide in the abstract; the inventory tracks whether specific decisions are actually being made, by whom, in time. Organizations with beautiful delegation matrices and terrible decision latency are common, because the matrix was never confronted with a dated queue. Run together, the queue's evidence continually corrects the matrix: owners who never appear can be pruned, and boundaries that generate repeated authority gaps get redrawn.

The instrument's scope needs a floor. Not every choice belongs in it — teams make hundreds of decisions a week inside their own authority, and an inventory that swallows them becomes a bureaucratic tax that discredits the idea. The entry test is consequence and span: choices that cross team boundaries, bind constrained capacity, set precedent, or carry material irreversibility. A portfolio typically holds dozens of those at a time. The inventory that governs them fits on a few pages.

## Starting the inventory by Friday

The inventory starts small and earns its place:

- Begin with decisions that affect multiple investments or constrained roles.
- Add the inventory beside schedule and risk reporting, not as another disconnected log.
- Review queue age, missed latest-useful dates, repeated owners, and decisions reopened because records were weak.
- Use recurring patterns to change delegations, evidence standards, or forum design.

The first population is archaeological: ask each program lead for the three questions they are currently waiting on, frame those properly, and the initial inventory exists by Friday. Expect it to be uncomfortable — several entries will turn out to have been open for months, and a few will have quietly passed their latest useful dates already. That discomfort is the instrument calibrating the organization's sense of what delay actually costs.

Repairing escalation culture is part of the rollout, and it is mostly modeling. The first few times a team escalates against an approaching latest useful date, the response teaches everyone watching: if the forum thanks the team and decides, dated escalation becomes normal practice; if the forum treats it as noise or the team's leader hears about it sideways, the inventory will fill with politely undated entries within a quarter. Leaders who want the instrument to work should engineer two or three early, visible, well-received escalations — the culture is set by cases, not by the rollout memo.

Cadence integration matters more than tooling. The inventory reviews inside the existing portfolio forum — ten minutes on aging items and approaching dates — rather than in a new meeting. The PMO curates framing quality and data honesty; owners decide; the forum handles what exceeds individual authority. A dedicated decision-review board is the failure mode, not the goal: it would recreate the friction the instrument exists to remove, one venue further from the work.

## What the record starts to show

A few quarters of inventory data answer questions the organization could previously only argue about:

- Median days from decision framed to decision closed, by forum and by subject.
- Decisions resolved inside their latest useful dates, and the trend.

- Queue age distribution, and which owners or forums hold the aging tail.
- Recurring subjects converted into standing policy, per quarter.
- Rework and idle capacity attributed to late answers, priced where possible.
- Decisions reopened, and whether the reopening matched a recorded condition.

The record has a quiet second life as institutional memory. New leaders inherit portfolios mid-story; a year of closed decisions with rationale attached is the fastest orientation an organization can offer — what was chosen, why, under what conditions it reopens. Teams stop relitigating settled questions because the settlement is findable. The inventory begins as a friction instrument and matures into the organization's decision history, which may be the more valuable asset.

The pricing measure deserves persistence even though it is the hardest. When a team can say 'the five-week policy wait cost eleven person-weeks of speculative rework,' decision friction stops being an abstraction in a governance deck and becomes a line item leaders recognize. Two or three priced cases per year sustain the instrument's mandate better than any adoption metric.

## The honest limits

Speed is not the sole objective. Some choices deserve deliberation, challenge, consultation, or legal review. The control distinguishes necessary decision work from unmanaged waiting and makes the consequence of elapsed time visible.

Nor is the instrument immune to gaming. Latest useful dates can be inflated to force attention, or framed conservatively to hide slack; owners can park items as 'awaiting evidence' indefinitely. The countermeasure is the record itself: dates that repeatedly prove wrong in the same direction, and evidence requests that never eliminate an alternative, are visible patterns the PMO reads back to the forum. An inventory audited against outcomes disciplines its own inputs within a few cycles.

The inventory also cannot supply courage. An organization whose executives will not choose between two sponsors' priorities will find its inventory full of beautifully framed, perfectly dated, unmade decisions — and that visibility is itself the contribution. The instrument converts a diffuse cultural complaint into a specific, attributable pattern: this forum, these subjects, this many days past useful. What leadership does with that evidence is, fittingly, a decision.

## The shift

Portfolios govern work; the proposal is to govern the choices that control the work with equal discipline. Named, owned, evidenced, dated, connected, closed — the same rigor schedules receive, applied to the questions schedules depend on.

Start where the pain already is. Every portfolio has two or three programs whose status reports have said 'awaiting direction' in some dialect for months. Framing those first — with their real owners, real dates, and real downstream costs — gives the instrument its opening proof, and gives those programs their answers.

The payoff arrives as absence: fewer mysterious slips, fewer retrospectives that blame estimation for what was actually an unaddressed question, fewer options lost to the calendar. Teams feel it as a change in the weather — questions asked in week one get answered while week one's options are still alive. That is what it means for an organization to be decisive: not that its leaders decide quickly, but that its decisions arrive while they still matter.

# Notes and sources

[1] Project Management Institute, "Pulse of the Profession 2026: Driving Success in Complex Projects," May 2026. https://www.pmi.org/learning/thought-leadership/driving-success-in-complex-projects

[2] CIO, "IT Governance: Best Practices for Aligning IT and Business Strategy." https://www.cio.com/article/272051/it-governance-a-formal-way-to-align-it-and-business-strategy.html

### About the author

Marco Policani is an enterprise portfolio, PMO, and AI operating-governance leader. He builds portfolio governance systems where AI carries the analytical load and named humans own every decision — an approach documented across case studies, walkthroughs, and working governance guides on his portfolio site.

Portfolio & governance library: policani.net/governance · Full portfolio: policani.net · LinkedIn: linkedin.com/in/marcpolicani