10

WHITE PAPER · PORTFOLIO & DELIVERY GOVERNANCE

A Portfolio Needs Scenarios, Not a Single Annual Plan

Pre-agree contingent choices before surprise forces them

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

July 2026

EXECUTIVE SUMMARY

Scenarios create options before a capacity change, dependency slip, or new demand forces an unprepared choice.

Single plans fail by preserving their own appearance: every initiative stays alive, delay spreads as fairness, and reprioritization happens in emergencies under sunk-cost pressure. This paper builds the scenario steering cycle — feasible alternative portfolio shapes, observable triggers with named owners, and pre-agreed first moves that are reversible by design.

01 Compose two genuinely feasible alternative shapes; a downside case that keeps every initiative is the baseline wearing a worried expression.

02 Set triggers by lead time and give each one an individual owner, so noticing is never the bottleneck.

03 Measure the cycle by emergency reprioritizations avoided — surprise should arrive as a convened decision with prepared options.

The argument

An annual plan is a useful commitment baseline and a poor description of every future condition. A small set of feasible scenarios lets leaders pre-agree what they will protect, defer, reshape, or stop when capacity, demand, dependencies, or assumptions change.

I learned to distrust single plans in the year a portfolio I supported absorbed three surprises before June: a hiring freeze in February, an acquired product line in April, and a platform vendor pushing a migration nobody had scheduled. None of the surprises was exotic. Any planning session could have imagined all three. What the organization had instead was one approved plan, so each surprise triggered the same ritual — an emergency reprioritization, conducted in a hurry, by whoever could attend, under exactly the conditions that produce bad portfolio decisions: incomplete information, sunk-cost pressure, and a room full of people defending their own initiatives.

The pattern to notice is not that plans meet surprises. It is what happens next. Most portfolios respond to a broken planning premise by preserving the plan's appearance: every initiative stays alive, delay spreads evenly, overload is distributed as a kind of fairness, and the plan document remains formally in force long after it stopped describing anything. The alternative response — choose again, deliberately, from prepared options — requires the options to have been prepared. That preparation is what scenario discipline buys.

Scenarios here are not forecasting exercises or strategic narratives. They are operating options: alternative portfolio compositions, tested for feasibility in advance, with the conditions that would trigger each and the moves each implies. This paper — written for portfolio executives, EPMO leaders, strategy leaders, and functional executives — lays out a six-part steering cycle for building and running them. PMI's complex-project research and Flyvbjerg's front-end due-diligence work both argue for confronting uncertainty before commitment rather than after [1][2]; the cycle itself was assembled the hard way, one absorbed surprise at a time.

What a single plan actually costs

The single plan's defects are timing defects. When a surprise arrives, the organization must simultaneously recognize the change, invent its response, negotiate that response across every affected party, and execute it — all while delivery continues burning capacity against the old assumptions. Each step has a natural delay, and the delays multiply. By the time the reprioritization lands, months of constrained capacity have gone to work the new reality had already demoted.

Negotiating under fire also selects for the wrong outcomes. In an emergency reprioritization, the initiatives that survive are the ones with the most effective advocates in the room that week — not the ones the portfolio would choose with a clear head. Sunk cost argues loudly, optionality has no constituency, and the sacrifice lands wherever resistance is weakest. Decisions of this kind are individually defensible and collectively incoherent, which is why portfolios that live this way feel permanently busy and strangely directionless at the same time.

The annual plan's defenders are right about one thing, and the scenario cycle should concede it clearly: organizations need a commitment baseline. Budgets, hiring, vendor contracts, and cross-functional promises all require a stated shape of record. The argument is not against having a plan; it is against having only a plan — a single composed future with no prepared neighbors, defended past its premises because the alternative to defending it is improvising in public.

There is a quieter cost as well: threats and opportunities that are never even recognized as portfolio events. Without named uncertainties and watch signals, a slow deterioration — a dependency slipping monthly, demand softening quarter by quarter — never crosses a threshold anyone owns. The portfolio does not decide to ignore it; the portfolio simply has no mechanism by which it could notice.

The scenario steering cycle

The cycle has six parts: choose the uncertainties that matter, build feasible alternative shapes, make each option's consequences visible, define observable triggers, pre-agree the contingent moves, and review assumptions on a cadence. It runs continuously beside the baseline plan — the plan remains the commitment of record; the scenarios are the prepared answers to the question the plan cannot ask.

EXHIBIT 1

Six moves convert a plan into prepared options

1. **Choose decision-relevant uncertainties**
   A scenario should vary conditions that would change portfolio choices, not catalog every imaginable event.
2. **Build feasible portfolio shapes**
   Each scenario must fit the capacity and dependency conditions it assumes.
3. **Show what each option protects**
   Leaders need to see the strategic and operating consequences of choosing a scenario.
4. **Define observable triggers**
   A scenario becomes operational when movement toward it can be recognized early enough to act.

5. **Pre-agree contingent moves**
   The portfolio should know which actions become candidates under each condition.
6. **Review assumptions, not theater**
   Scenario review should ask what changed and whether the active portfolio still fits, not repeat the annual prioritization ritual.

Choose decision-relevant uncertainties

Scenario work fails most often at the first step, by choosing uncertainties that are interesting rather than decisive. A scenario should vary a condition that would change actual portfolio choices — sequencing, funding, scope, staffing, dependency posture. If leaders would do the same things in both branches, the uncertainty is background, not structure.

Signal sourcing belongs to this step as well: an uncertainty is only usable if some observable data speaks to it. Capacity uncertainties read from attrition and requisition aging; dependency uncertainties from delivery performance against commitments; demand uncertainties from pipeline, volumes, or the regulator's published calendar. If no observable proxy exists, the uncertainty can still shape thinking, but it cannot carry a trigger — and the set should favor the watchable.

The practical filter is a two-question test applied to each candidate: could this move materially within the planning horizon, and if it did, which commitments would we want to change? Uncertainties that pass both — a possible regulatory date, the retention of a critical capability, a dependency's delivery confidence, a demand assumption with real variance — become scenario axes. Two or three are enough. Beyond that, the combinatorics manufacture work without adding decisions.

Keeping the set small also forces an honest ranking conversation: of everything that could change, what would actually move this portfolio? I have watched leadership teams discover, mid-argument, that they disagreed about the answer — which meant they had been running the portfolio on unexamined, incompatible assumptions all along. Surfacing that disagreement in a planning room, before anything has broken, is scenario discipline already paying for itself.

Build feasible portfolio shapes

Each scenario must be a portfolio that could actually run under its assumed conditions. This is the step that separates operating scenarios from slideware. A downside case that lowers the benefit forecast while leaving the same forty initiatives staffed by the same overcommitted people is not a scenario; it is the baseline plan wearing a worried expression.

Composition is the real work: under this scenario's constraints, which work continues at full shape, which narrows, which re-sequences, which stops, and which new work enters? The constrained roles, shared environments, and decision forums must all clear — if the reduced-capacity scenario still assumes the departed architect's output, it fails its own premise. Expect each shape to require genuine subtraction. A candidate scenario in which nothing stops is usually evidence that the composition was negotiated to avoid offense rather than built to fit the constraint.

Feasibility testing has a side benefit that repays the effort even if no trigger ever fires: it reveals which initiatives are load-bearing across every future — they appear in all shapes — and which are conditional passengers. That distinction alone sharpens the next intake round, because a proposal that would survive no adverse scenario is a proposal the portfolio is carrying on pure optimism.

Show what each option protects

A scenario earns a decision only when leaders can see what choosing it protects and what it gives up. Financial ranking compresses exactly the wrong things: resilience, regulatory exposure, customer commitments, capability building, and timing options all vanish into a net-present-value ordering that no executive actually believes.

The working format is a consequence view, one page per shape: outcomes protected, options preserved, risks accepted, capabilities built or abandoned, commitments sacrificed, and the assumptions the shape leans on hardest. Uncertainty stays visible — ranges rather than point estimates, and dissent noted where the team disagreed about a consequence. The goal is not a winner computed by spreadsheet; it is a choice whose costs the choosers have looked at directly.

This format changes the meeting in a way worth anticipating: it makes the sacrifice discussable before it is urgent. Executives who would resist cutting a program in an emergency — publicly, with a sponsor watching — can agree calmly that under scenario B, that program is the deferral. The agreement is cheaper by months when it is hypothetical, which is precisely the point of making it early.

Define observable triggers

A scenario without a trigger is a document; with one, it is a standing decision waiting for its condition. The trigger discipline: for each uncertainty, a small number of observable signals, a threshold range that means 'convene the choice,' a named owner who watches, and a stated observation cadence. Observable is the operative word — 'market deterioration' is a mood, while 'two consecutive quarters of pipeline below X' is a fact someone can own.

Trigger ownership is individual by design. A threshold watched by a team is watched by no one, because each member reasonably assumes another is watching. The named owner's obligation is small — observe on cadence, log the value, raise the flag at the threshold — and precisely because it is small, it survives busy quarters that would kill a heavier duty.

The design tension is false alarms against late alarms, and it should be settled by lead time: the threshold belongs wherever the contingent move still has room to execute well. A trigger that fires when the response is already forced was set for observers, not deciders. Where the signal is noisy, use a confirmation window rather than raising the threshold — react to sustained movement, not spikes, but react while reacting is still cheap.

Note what the trigger does and does not do: it convenes the decision, it does not make it. Leaders retain the judgment call; the trigger removes only the possibility of not noticing, and the political cost of being the person who raises the alarm. Institutionalizing the alarm is half of what makes contingent governance function in real organizations.

Pre-agree contingent moves

When a trigger fires, the difference between a prepared portfolio and an unprepared one is whether the next meeting starts from options or from scratch. Pre-agreed moves are the options: for each scenario, the candidate actions — defer named work, narrow a scope, protect a constraint, switch a sequence, release a dependency, accelerate a capability — with the decision owner, the affected commitments, the first reversible step, and the communications each move implies.

The first reversible step deserves emphasis, because it is what lets the portfolio respond before certainty arrives. Committing fully to a scenario on a trigger's first firing over-steers; doing nothing until the scenario is undeniable wastes the preparation. The prepared middle: execute the smallest move that preserves options while evidence accumulates — pause the next tranche rather than cancel, hold the architect rather than reassign, open the vendor conversation rather than sign. Each buys information at low cost.

Pre-agreement also has to survive its own success. The conditions for returning to the prior shape belong in the move design — otherwise every contingent action becomes a one-way door by default, and the portfolio ratchets toward conservatism with each firing. A move that can be entered, evaluated, and exited is a steering instrument; a move that can only be entered is a slow-motion emergency.

Review assumptions, not theater

The cycle stays alive through a short recurring review with a narrow agenda: what moved among the signals, which assumptions changed, whether the active shapes remain feasible, and which scenarios have expired. Expiry is a health signal, not a failure — a scenario built around a resolved uncertainty has done its work and should leave the set, replaced by whatever uncertainty has since earned a seat.

The review's enemy is its own institutional gravity. Left unattended, scenario review drifts toward annual-planning theater: full re-prioritization, fresh narratives, decks that grow each cycle. The countermeasure is scope discipline — the review examines assumptions and feasibility, and only a fired trigger or an expired premise opens composition questions. Twenty minutes on the quarterly portfolio agenda sustains the discipline; a standalone scenario summit kills it within a year.

Review is also where the record accrues. Signals observed, assumptions revised, triggers approached or fired, moves executed and unwound — a few quarters of this history shows the organization how its uncertainty actually behaves, which recalibrates both the thresholds and the leadership instinct for what deserves preparation. Portfolios get measurably better at scenario choice by doing it, but only if the doing leaves a trace.

One uncertainty, run through the cycle

A concrete chain shows the parts meshing. The uncertainty: a portfolio's largest program depends on a vendor platform whose last two releases slipped, and the vendor has announced a restructuring. The two-question test passes — the risk could mature this year, and it would change sequencing, staffing, and a contract decision.

The shapes: the baseline assumes the vendor delivers within one quarter of promise. Shape B assumes a two-quarter slip — the dependent program re-sequences its integration phase behind a workstream that needs no vendor input, a bridging capability is scoped for the interim, and the contract's renewal option converts to a shorter term. Shape C assumes vendor failure — migration to the alternative platform begins, two mid-tier initiatives defer to fund it, and the constrained integration team moves entirely to the migration. Both shapes clear the constraint model; both were composed calmly, in an afternoon workshop, months before anything happened.

The triggers: vendor release performance against committed dates, tracked by the program's vendor manager, with a threshold of one further slip beyond four weeks; and a qualitative signal — loss of the vendor's named engineering lead — that convenes review immediately. The moves: shape B's first step is reversible (hold the integration start, open the contract conversation); shape C's first step is a paid discovery engagement with the alternative platform, deliberately small.

When the vendor slipped again in August, the review convened within a week, executed shape B's first step, and set a confirmation window for shape C's discovery step. Total elapsed decision time: nine days, most of it scheduling. The unprepared version of that same event — I have lived it — consumes a quarter, and the migration begins from a standing start with the integration team already burned.

The cycle under pressure

The parts hold together through a specific dependency chain. Triggers are only as good as the uncertainties chosen; moves are only executable if the shapes were feasible; the review only means something if triggers are observable. When the cycle disappoints, trace backward: a trigger nobody acted on usually points to a move nobody actually pre-agreed, which points to a consequence view nobody confronted, which points, most often, to shapes that were never honestly composed.

EXHIBIT 2

The contingent choice each part prepares

Choose decision-relevant uncertainties
Exclude uncertainty that would not alter a portfolio commitment.

Build feasible portfolio shapes
Reject scenarios that preserve every priority by hiding overload.

Show what each option protects
Choose the portfolio shape whose tradeoffs leaders are prepared to own.

Define observable triggers
Convene the choice when the trigger enters its action range, not after the old plan has failed.

Pre-agree contingent moves
Execute the smallest move that preserves options while confirming the scenario.

Review assumptions, not theater
Update the set when the decision landscape changes rather than preserving scenarios for presentation continuity.

And the honest test of the whole apparatus is behavioral, applied at the first real firing: did the portfolio's response start from the prepared options? If the answer is that leaders convened, opened the scenario file, and negotiated from its moves — the cycle works, whatever its imperfections. If the emergency ritual reappeared with the scenario file unopened, the organization has scenario documents, not scenario governance, and should either fix the pre-agreement step or stop paying for the documents.

Running the cycle beside the annual plan

The cycle installs incrementally, and modestly:

- Build a baseline plus two materially different feasible alternatives, testing feasibility against constrained roles and dependencies before debating preference.
- Limit triggers to signals leaders can observe and act on.
- Run a short scenario review before quarterly portfolio decisions and after material shocks.

Sequence matters: compose the alternative shapes first and let the consequence views generate the trigger discussion, not the reverse. Teams that start with trigger brainstorming produce elaborate sensing for options they never build. Teams that start with two honest shapes discover their triggers naturally — the act of composing 'what we would run if capacity fell' surfaces exactly which capacity signal matters and where its threshold sits.

Funding mechanics deserve early attention, because contingent moves need contingent money. A pre-agreed move that requires a fresh business case at firing time is not pre-agreed; it is pre-discussed. The workable patterns are a small contingency tranche held at portfolio level with drawdown authority tied to named triggers, or pre-approved reallocation authority — the forum may move up to a stated amount between named initiatives when a trigger fires, with retrospective review. Either pattern converts the first reversible step from a budgeting exercise into an operating one.

Sponsor politics also need tending in advance. The owner of the initiative that defers under shape B should not discover that fact at firing time — the consequence view conversation is where that sponsor hears it, argues it, and either changes the composition or accepts it. A sacrifice negotiated hypothetically is a governance agreement; the same sacrifice imposed during a crisis is an ambush, and ambushed sponsors litigate. Much of the cycle's value is simply moving those negotiations to the cheap end of the timeline.

Ownership follows the usual portfolio split. The EPMO or portfolio function curates the cycle: maintains the shapes' feasibility, watches the signal owners watch, keeps the review short. The choices — which shape, when to move, when to return — belong to the executive forum. The one new role is the trigger owner — a named individual per signal, for the reasons already given.

The cycle also has a natural relationship with rolling forecasts and quarterly funding models, for organizations that use them. Rolling processes re-examine the plan on a calendar; the scenario cycle re-examines it on evidence. They complement rather than compete — the quarterly re-forecast is a scheduled review, and a fired trigger is an unscheduled one, and both feed the same forum. What the scenario cycle adds that rolling planning lacks is the prepared option: re-forecasting tells you the premise broke, but only a composed shape tells you what to do about it by Friday.

What the record should show

A working cycle leaves evidence:

- Uncertainties with named signals, owners, and current values — reviewed on cadence.
- Alternative shapes that remain feasible against the current constraint model.
- Triggers that fired and convened a decision within their intended lead time.
- First moves executed reversibly, then confirmed, extended, or unwound.
- Scenarios retired on resolved uncertainty and replaced by live ones.
- Emergency reprioritizations per year — the number the cycle exists to reduce.

Read the record annually for calibration as well as accountability: triggers that fired with no consequential movement argue for higher thresholds; deteriorations that arrived unannounced argue for new signals; moves that were never exercised in any firing may be decoration. The cycle is itself an assumption set, and it deserves the same periodic disconfirmation it imposes on the plan.

The last line is the outcome measure. Scenario governance succeeds exactly to the degree that the organization stops holding emergency reprioritizations, because the events that would have been emergencies arrive instead as convened decisions with prepared options. One avoided emergency typically repays a year of review minutes.

The strongest counterargument

Scenarios do not make forecasts more certain and should not become disguised predictions. Their value is preparedness: they expose consequences, preserve options, and shorten the time between recognizing a changed condition and making a coherent portfolio choice.

A second objection has more force than it is usually given: scenario work can become its own industry, consuming planning capacity that delivery needed. The defense is the cycle's deliberate smallness — two or three uncertainties, two alternative shapes, triggers on one page, review in twenty minutes. An organization that cannot keep scenario work this small should suspect that the exercise has become intellectually interesting, which is a different thing from operationally useful. The apparatus deserves exactly as much investment as the emergencies it prevents, and no more.

The shift

A single plan asks one future to arrive on schedule and treats every departure as an exception to be absorbed. A scenario cycle assumes the future will negotiate — and prepares the portfolio's side of the negotiation in advance: uncertainties named, shapes composed, consequences faced, triggers owned, first moves reversible, return paths designed.

Opportunity deserves the same machinery, a point easily lost because scenario language skews defensive. An acquisition target becoming available, a competitor stumbling, a regulation opening a market — these are triggers too, and a portfolio with a pre-agreed 'upside shape' can commit capacity to a sudden opening while competitors are still scheduling their emergency meeting. Preparedness is symmetric; only the habit of treating it as insurance makes it feel otherwise.

The deeper change is in what surprise means. In a single-plan portfolio, surprise is a failure of planning and arrives as crisis. In a scenario-governed portfolio, most surprise is a trigger doing its job — recognized early, met with prepared options, decided while deciding is still cheap. The plans look similar on paper. The difference appears the week something breaks.

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] Bent Flyvbjerg, "Quality Control and Due Diligence in Project Management," 2013. https://arxiv.org/abs/1302.2544

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