10

WHITE PAPER · PORTFOLIO & DELIVERY GOVERNANCE

Stage Gates Work Only When They Can Kill Projects

Decision rights that protect scarce capacity

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

July 2026

EXECUTIVE SUMMARY

A gate earns its place only when it can continue, redirect, pause, or stop the next commitment.

A gate that cannot stop anything launders weak commitments into approved ones, because each passage stamps the work with the appearance of scrutiny. This paper rebuilds the gate as a commitment decision: delegated authority, evidence proportioned to the next tranche, five credible outcomes, independent evidence owners, and conditions that expire.

01 Audit the gate system behaviorally: in the last year, what did gates actually change about money, capacity, scope, or reliance?

02 Derive evidence demands from the next commitment, not a standing checklist — a thick pack is often camouflage for one unresolved question.

03 Attach trigger thresholds and latest-valid dates to every approval so it stops functioning as permanent permission.

The argument

A stage gate governs only when it connects evidence to a consequential choice. If every project advances regardless of what the review finds, the gate is a status meeting wearing the language of control.

The most revealing question I ask about any gate process is not what the templates contain. It is: when did this forum last stop something? In one portfolio I reviewed, the answer was never — four years of gate reviews, hundreds of approvals, zero terminations, zero redirections. Leaders around that table could request clarification, express concern, and ask for another slide. What they could not do, in any practical sense, was decline. Staffing was committed, vendor contracts were signed, and the executive sponsor had announced the initiative at a town hall. The gate had an agenda, a quorum, and no authority.

That arrangement rewards momentum over evidence. Early uncertainty, which is normal and cheap to act on, converts steadily into sunk cost, which is expensive and politically protected. By the time the weakness is undeniable, stopping costs more, embarrasses more people, and lands on a successor. A gate that cannot kill a project does not merely fail to help; it actively launders weak commitments into approved ones, because each passage stamps the work with the appearance of scrutiny.

This paper is written for portfolio boards, sponsors, PMO leaders, and investment owners. It describes a gate design built on five properties: delegated decision authority, evidence proportioned to the next commitment, five credible outcomes, evidence owners separated from advocates, and conditions that expire. Stage-gate practice has a long commercial pedigree [1], and the front-end due-diligence case has strong academic support [2]; the design here is drawn from my own portfolio and PMO work, and it is a management instrument rather than a guarantee.

How gates decay into ceremony

Nobody designs a toothless gate. Gates decay, and the decay follows a pattern. First, the forum's decision rights are left implicit at setup, usually to avoid an awkward conversation about who can overrule whom. Second, the gate pack grows: each past incident adds a section, until preparation is a project of its own and the review skims a hundred pages to discuss three. Third, outcomes narrow in practice to pass and pass-with-comments, because the first time a stop was plausible, someone discovered the contracts were already signed.

The end state is familiar. Teams prepare for weeks, present for an hour, and receive conditions that no one tracks. The real decisions — whether to fund, whether to staff, whether to rely — happen elsewhere, earlier, in conversations the gate never sees. The forum has become a theater with excellent production values, and its main output is the retrospective legitimacy it confers on choices already made.

The decay is rarely visible from inside, because every step is locally reasonable. Leaving decision rights vague avoided a real conflict. Each pack section answered a real incident. Each pass-with-comments reflected genuine sympathy for a team that had worked hard. The system degrades through accumulated courtesy, which is why fixing it requires design rather than exhortation — asking a courteous forum to be tougher changes one meeting, not the structure that made courtesy the path of least resistance.

The repair is not more rigor in the pack. It is restoring the connection between what the review learns and what the organization does next. Everything in this paper serves that single connection.

A gate that earns the next commitment

The design principle: a gate is the moment the organization decides whether to extend, reshape, or withdraw its next commitment — money, scarce people, dependency, or reliance. Five properties make that decision real rather than ceremonial. They stand together; the later sections show what each looks like when present, how it fails when absent, and what evidence tells the difference.

EXHIBIT 1

Five properties separate a gate from a status meeting

1. **Delegate real decision authority**
   The gate forum needs explicit authority to continue, redirect, narrow, pause, or stop the next commitment.
2. **Ask for evidence proportionate to commitment**
   Evidence should answer the risk created by the next tranche, not satisfy the same checklist at every stage.
3. **Keep all five outcomes credible**
   Advance is one valid outcome among several; redirection, narrowing, pausing, and stopping are evidence of governance working.
4. **Separate evidence owners from advocates**
   The sponsor may advocate for the investment, but critical evidence needs an owner accountable for its integrity.

5 **Make conditions expire**

A gate decision is based on evidence that can age; approval should identify the assumptions that would reopen it.

Delegate real decision authority

The forum needs explicit authority to continue, redirect, narrow, pause, or stop the next commitment. Not influence — authority. Before any evidence is prepared, the gate's charter should publish its decision rights, quorum, reserved matters, and escalation path, and the resource owners should have accepted that the forum's decision moves their money and people.

The absent version is more common: participants advise while the actual decision sits with a sponsor, a CFO, or an executive committee that meets quarterly. The project team, reasonably, treats discussion as approval and keeps moving. When a later crisis triggers the question 'who approved this?', the answer is a diffusion: the gate assumed the sponsor decided, the sponsor assumed the gate did, and the minutes record only that concerns were noted.

The test for real authority is behavioral, and it is worth running deliberately: pick a low-stakes gate and have the forum attach a genuine condition — a scope reduction, a staffing change, a delayed tranche. If the condition takes effect without renegotiation, the authority is real. If the project absorbs the condition into a slide and continues unchanged, redesign the forum before trusting it with anything that matters. A forum that cannot move money, capacity, scope, or reliance should not be maintained as if it could; either give it the authority or stop calling it a gate.

Authority also needs a boundary map, because gates share the decision space with other forums. Architecture boards, security reviews, funding committees, and executive sponsors all hold pieces of the commitment. When the pieces overlap without a map, projects venue-shop: the answer refused at the gate is sought from the sponsor, and the sponsor's yes quietly outranks the forum's no. The charter should therefore state not only what the gate decides but what it explicitly does not — and what happens when two forums disagree. An escalation path that has never been exercised is a hypothesis; run it once on purpose before it is needed in anger.

Ask for evidence proportionate to commitment

The gate's information demand should be derived from the next tranche, not from a standing checklist. State the commitment first: the next six months means this spend, these named people, this new dependency. Then ask what must be true for that commitment to be sound, and what evidence would make a skeptical decision maker comfortable extending it. That derivation produces a short, specific evidence map — each material premise linked to a test, a result, an owner, and the uncertainty that remains.

Checklists produce the opposite: large packs of low-value material in which the decisive uncertainty hides. A team can satisfy forty template sections while the single question that should decide the gate — will customers adopt this, can the integration actually be built, does operations have capacity to absorb it — remains exactly as open as it was at the previous gate. Volume is not diligence. In practice volume is often camouflage, and experienced forums learn to read a thick pack as a signal that something specific is being buried.

Proportionality cuts both ways. Reversible exploration should advance on modest evidence; demanding production-grade proof of a six-week experiment teaches teams to route around the process. Irreversible spend, enterprise dependency, or reliance by other programs deserves the strong form: independent checks, tested premises, and explicit statements of what is still unknown. A gate that asks the same questions at every stage is asking the wrong questions at most of them.

A practical format that holds the discipline: a one-page evidence map on top of the pack, with everything else as reference. The page lists the next commitment in one sentence, the three to five premises it depends on, and for each premise the test performed, the result, the owner, and what remains unknown. Forums that adopt the page find their reviews change shape within two cycles — the conversation moves from slide sequence to premise quality, and preparation effort drops because teams stop producing material nobody asked a question about.

Keep all five outcomes credible

Advance is one valid outcome of a gate. The other four — redirect, narrow, pause, stop — must remain live options in fact, not merely on the process diagram. Credibility is empirical: teams believe the outcomes they have seen. A portfolio in which no one can recall a stop decision is a portfolio in which every team prepares for exactly one outcome, and prepares accordingly — defending the original plan rather than revealing what has changed.

Making the other outcomes real requires defining their consequences in advance. What happens to the released capacity after a stop? Where does a paused project's knowledge go, and what conditions bring it back? A narrow decision changes which commitments, held by whom? If those questions have no standing answers, each non-advance outcome becomes a bespoke negotiation, and forums default to advancing because it is the only outcome with a worked path.

The cultural half is protecting the messenger. The team that surfaces disconfirming evidence at gate three has saved the organization the remaining tranches; if the institutional response is to punish the project's leadership, the next team will bury the signal until it becomes undeniable — at several times the cost. A practical marker: when a stop decision is communicated, the language should credit the evidence, distinguish a sound stop from a delivery failure, and be visibly reflected in how the affected people are redeployed. Teams watch what happens to the last team.

Redirect and narrow deserve particular attention, because they are the workhorses. Full stops are rare in a decently filtered portfolio; what evidence more often shows is that the original shape is wrong — the market case holds for one segment, the platform bet holds without the speculative second phase, the rollout survives contact with operations only at half speed. A forum fluent in reshaping decisions extracts value from weak projects that a pass/kill binary would either wave through whole or destroy whole. The gate's vocabulary should be as granular as the evidence.

Separate evidence owners from advocates

A sponsor is entitled to advocate. The gate's problem begins when the same people whose funding depends on approval are the sole authors of the forecasts, readiness claims, and benefit assumptions the forum relies on. That is not a character flaw; it is a structural conflict that predictably shades estimates, and it is why gate evidence needs owners who are accountable for its integrity rather than its persuasiveness.

The practical mechanics are modest. High-consequence premises get independent or cross-functional challenge: finance re-derives the benefit math, architecture reviews the integration claim, operations confirms the absorption assumption. Reference-class comparisons — what did the last five projects of this shape actually cost and deliver — counter the optimism that inside views generate [2]. Assumptions surface in the body of the pack with sensitivity ranges, not in an appendix. Dissent is recorded, because a forum that hears only agreement is hearing a performance.

When the forum cannot distinguish advocacy from evidence on a premise that matters, it has two honest moves: commission the missing test, or shrink the commitment to what the untested premise can safely carry. Approving the full tranche on advocacy alone is the quiet version of skipping the gate entirely.

One portfolio pattern is worth naming: the same premise, unchallenged, across many packs. Benefit ramp curves that always reach plan in month seven. Adoption assumptions copied between business cases. Integration estimates that never include the system everyone knows is fragile. Independent evidence ownership at the gate catches the instance; the PMO, reading across gates, catches the pattern — and a corrected pattern repriced every future business case at once, which is more value than any single gate decision delivers.

Make conditions expire

A gate decision is a judgment about evidence at a moment. The evidence ages: costs move, sponsors change, dependencies slip, market assumptions drift. A six-month-old approval can describe a project that no longer exists, yet in most portfolios the approval functions as permanent permission — the work continues because it was once approved, not because the approval's basis still holds.

The control is to attach expiry to the conditions that matter. Each material assumption behind an approval gets a trigger threshold and a latest-valid date, with a named owner responsible for bringing the work back to the forum when either is crossed. The trigger list is short by design — the three or four premises the approval actually leaned on — and it converts 'we should keep an eye on that' into a mechanism that reopens the decision while options still exist.

This is the property that makes the other four durable. Delegated authority, proportionate evidence, credible outcomes, and independent challenge all operate at the gate. Expiring conditions extend their reach between gates, which is where most of a project's life happens and where most of its assumptions break.

A worked example of the mechanism: an approval rests on a vendor's committed integration date and a benefits case priced against current volumes. The gate attaches two triggers — vendor slip beyond six weeks, volume forecast revised more than ten percent — and a latest-valid date matching the next tranche. Four months later the vendor slips nine weeks. Without the trigger, that fact would surface as a schedule variance in a status report, absorbed and re-baselined. With it, the work returns to the forum while the tranche is still unspent, and the choice — wait, re-sequence, or proceed with a contained scope — is made with options intact.

Five properties, one failure mode

The properties reinforce one another, and they share a failure mode: each can be simulated. A charter can list decision rights the forum never exercises. An evidence map can be produced by the advocates. A stop can exist in the process manual and nowhere else. Conditions can be attached and never tracked. The simulation looks like governance from every angle except the one that matters — whether the portfolio's actual commitments moved in response to what reviews found.

EXHIBIT 2

The decision each property protects

Delegate real decision authority
If the forum cannot change money, capacity, scope, or reliance, redesign the forum rather than adding ceremony.

Ask for evidence proportionate to commitment
Advance reversible exploration with modest evidence; demand stronger proof before irreversible spend, dependency, or enterprise reliance.

Keep all five outcomes credible
Choose the path that preserves the most value and optionality, not the path that best protects the original narrative.

Separate evidence owners from advocates
Commission a test or reduce the commitment when the forum cannot distinguish advocacy from evidence.

Make conditions expire
Reopen, narrow, or pause when the basis for approval no longer describes the work.

That is why the audit question for a gate system is behavioral, and blunt: in the last year, what did gates change? Money moved, tranches narrowed, projects redirected, work stopped, conditions enforced. If the honest answer is 'nothing material,' the system is ceremonial regardless of its documentation, and the redesign should start with authority — the property whose absence makes the other four decorative.

The gate system as a portfolio instrument

Individual gates decide individual commitments; the system of gates manages a scarcer resource — senior decision attention. A portfolio running forty projects through five-stage governance is implicitly scheduling two hundred consequential reviews over their lifecycles, most of them past the same dozen executives. That arithmetic explains most gate theater: overloaded reviewers skim, skimming rewards polish over substance, and polish is what advocacy produces best.

Two design responses keep the arithmetic honest. Proportionate routing: low-consequence, reversible work moves on a consent agenda — evidence filed, objections invited, no meeting — reserving live review for commitments that are large, irreversible, or contested. And batching by decision type rather than by project: a quarterly session that decides five tranche extensions together forces the comparative question a single-project review never asks — which of these five deserves the constrained architects — and comparison is where portfolio governance actually lives.

Restoring teeth to an existing gate process

The repair is mostly subtraction and rewiring rather than new machinery:

- Rewrite every gate agenda around the decision required, not the presentation sequence.
- Limit the core pack to the evidence that could change that decision.
- Record outcome, conditions, dissent, owner, and expiry in the same meeting.
- Audit whether capacity and funding actually moved after redirect, pause, or stop decisions.

Sequence the change at one gate first — ideally the tranche decision on a mid-sized investment, where stakes are real but the politics are survivable. A visible redirect or narrow decision at that gate, executed cleanly, will do more to reform the portfolio's expectations than a revised process manual distributed to everyone. Then extend the design outward, letting the early cases supply the worked examples the manual lacks.

What the first visible case can look like: a platform investment arrives at its tranche gate with adoption running at a third of the case. Under the old system, that is a yellow slide and a commitment to a recovery plan. Under the new one, the forum narrows — the next tranche funds the two adopted modules, the remaining scope returns to intake as a fresh proposal, and the freed capacity moves the same month to a program that had been waiting on those engineers. No drama, one page of record, and every sponsor in the portfolio recalibrates what a gate means.

The PMO's role is custodial, not judicial: maintain the method, prepare the evidence maps, track conditions to their expiry dates, and surface cross-portfolio patterns. The decisions belong to the forum, and the accountability for each investment stays with its sponsor. A PMO that starts making gate decisions has not gained authority; it has absorbed blame that belongs elsewhere.

Expect the second and third reviews under the new design to be the hard ones. The first runs on novelty; by the third, the deferred conflicts arrive — a sponsor who escalates around a narrow decision, a resource owner who quietly restaffs a paused project. Each of those is the old system testing the new one, and each holds or folds on whether the published escalation path and the resource-movement audit actually operate. Treat the first contested decision as the real go-live. Everything before it was rehearsal.

One caution on scope: do not convert every recurring meeting into a gate. The design in this paper is for commitment decisions — moments where money, scarce people, dependency, or reliance changes hands. Applying kill-capable review weight to routine delivery checkpoints burns exactly the decision attention the system exists to conserve, and teaches teams that governance is weather rather than signal.

Evidence the system has teeth

A working gate system leaves marks in the portfolio record:

- Non-advance outcomes actually used — redirects, narrowings, pauses, stops — and the capacity they released.
- Where released capacity went, and how quickly it was productive elsewhere.
- Conditions attached versus conditions tracked to resolution or expiry.
- Premises independently checked, and the frequency with which the check changed the estimate.

- Age of the evidence behind each active approval.
- Gate preparation effort relative to the size of the commitment being decided.

The last measure — preparation effort against commitment size — is the ceremony detector. When a team spends six person-weeks preparing to defend a two-person-month tranche, the system is consuming more than it governs. Healthy proportion runs the other way, and restoring it is usually the first visible benefit teams notice from the redesign.

None of these is a score, and one of them deserves a warning: a rising stop count is not automatically health. It can mean the gates are working, or that intake quality collapsed, or that the portfolio is churning through bad bets it should never have admitted. The measure earns its meaning from the recorded reasons. What the set reliably exposes is the ceremonial system — because a ceremonial system produces clean zeros all the way down.

The strongest counterargument

Strong gates can still become slow or political. The answer is not maximum documentation; it is a small number of consequential reviews, clear delegation, and evidence scaled to reversibility and exposure. Routine low-risk decisions should stay with the team.

The political objection is real as well: a forum that can stop work can be captured by whoever controls its membership, and a kill decision can be a weapon in a budget fight. The defenses are the same properties that make the gate work — published decision rights, evidence a skeptic can inspect, recorded dissent, and reference-class comparisons that anchor debate to observed history. A gate is harder to capture in daylight. Where the politics are genuinely toxic, the gate is not the problem, and no gate design will be the solution.

There is also a sharper objection: that kill-capable gates make teams risk-averse, since no one proposes ambitious work into a process that might stop it. The experience runs the other way when the design is right. What makes teams cautious is ambiguity about how they will be judged. A gate with explicit evidence standards, proportionate demands, and a protected route for bad news is safer for ambitious work than a ceremonial gate whose real decision criteria are political and hidden. Teams do not fear clear tests. They fear unwritten ones.

The shift

A gate that cannot change the commitment was never governing it. The shift this paper argues for is from reviewing projects to deciding commitments: every gate ends with the next tranche extended, reshaped, or withdrawn, on evidence a skeptic could inspect, under conditions that expire before they can quietly become permanent.

Portfolios that make this shift do not stop more work than they should. They stop it earlier — while the capacity still has somewhere better to go, and while stopping is a decision rather than an admission.

And the projects that pass gain something too: an approval that survived a forum with the power to refuse it. In a portfolio where gates have teeth, advancement is information. Where they do not, it is only momentum with a signature.

Notes and sources

[1] Stage-Gate International, "Gate Governance." https://www.stage-gate.com/solutions/services/gate-governance/

[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