WHITE PAPER · PORTFOLIO & DELIVERY GOVERNANCE

Why the Kill Gate Doesn't Fire

Let someone who did not approve the work make the call

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

September 2026

EXECUTIVE SUMMARY

Set a minimum result at approval, judge the sponsor on the case, and if the program falls short, let someone who did not approve the work decide.

A committee can hold full authority to stop a failing program and still extend it quarter after quarter. Tougher gates were the standard remedy. The escalation research shows why authority alone leaves the pattern in place: the sponsor who made the commitment argues the case before a committee that shared the approval, and a small set of studies found the pull to continue stronger when a responsible decision maker shares authority or is monitored. The design here moves the decision instead of replacing anyone.

01 People personally responsible for a commitment put more into it, and the effect was stronger when a responsible decision maker shared authority or was monitored.

02 Record the minimum result and its date at approval, before anyone has a record to defend, and keep it separate from the forecast.

03 Tell sponsors at approval they will be judged on the case and how early they raise a miss; judging on outcomes did not reduce commitment.

04 If the program misses its minimum, the sponsor recommends and a reviewer who did not approve the work decides; that reviewer can still continue it.

CONTENTS

Executive summary

A portfolio committee can hold full authority to stop a failing program and still extend it quarter after quarter. The standard remedy, set out by the founders of Stage-Gate Inc., is tougher Go/Kill gates, because projects that begin without serious scrutiny are rarely killed.

The escalation research explains why authority alone can leave the pattern in place. People personally responsible for a commitment put more resources into it when it goes badly, and projects that escalate end worse. A funding review puts the sponsor who made the commitment in front of a committee that shared the approval. A meta-analysis found the responsibility effect stronger when authority over the decision was shared or the responsible decision maker was monitored, although only four studies tested that condition.

Record the minimum result before any results arrive. Tell the sponsor at approval that the evaluation rests on the soundness of the case and how early a miss is raised. If the program falls short of it, the sponsor recommends and a reviewer who did not approve the work decides, and that reviewer can still continue it. The combined practices have not been tested together.

The quarter the gate extended

A finance-system consolidation is eighteen months into a three-year plan. Two of six business units have moved to the new platform, each migration has cost well above its estimate, and the next round of funding covers the largest unit. The gate charter lets this committee continue, narrow, pause, or stop the work, and nobody in the room disputes that authority.

The sponsor asks for one more quarter, names a new integration partner expected to cut migration cost, and points out that the largest unit has the cleanest data. The committee members approved the first two rounds of funding. They ask careful questions, request a revised migration schedule, and release the funding. A quarter later the conversation repeats over a larger number.

The scene is a composite. The committee had the authority to stop the program and did not use it. The usual diagnosis is that the gate needs sharper criteria or members with more backbone, and the usual remedy is to strengthen the gate.

Tougher gates were the first remedy

Robert Cooper and Scott Edgett, founders of Stage-Gate Inc., set out that remedy. Many projects, they write, move too far into development without serious scrutiny, and once a project begins "there is very little chance that it will ever be killed" [1]. Their prescription is tough Go/Kill decision points where managers decide whether to continue [1]. Portfolios that adopted it gave their committees the authority the opening scene shows.

The escalation research explains why authority and scrutiny can still leave failing work funded. Mark Keil, Joan Mann, and Arun Rai surveyed IS audit and control professionals about projects that did and did not escalate. Their results suggest that 30 to 40 percent of IS projects show some degree of escalation, and escalated projects had significantly worse perceived implementation and budget-and-schedule performance [2]. The construct that best separated escalated from non-escalated projects was the completion effect, which holds that the pull to finish grows as a project nears its end [2]. The predictors in that study describe the pull on the people deciding.

What a gate review does to the sponsor

The continue-or-stop decision at a gate goes to the people with the most at stake in the original choice. The sponsor wrote the business case and argued for it. The committee members approved the earlier funding rounds. Stopping the program requires them to accept that their earlier judgment was wrong, or that money already spent will not return what they promised.

Barry Staw's 1976 experiment isolated what that position does to a decision. In a simulated business investment, 240 business school students allocated funds and then learned the results. Those personally responsible for a negative outcome committed the greatest amount of additional resources to the course they had chosen [3]. A meta-analysis by Dustin Sleesman and colleagues, covering more than 35 years of escalation research, found personal responsibility for the initial decision positively related to escalation, as were sunk costs and proximity to project completion [4].

The same meta-analysis tested a condition that closely resembles a gate review. It counted responsibility as shared when more than one person held authority over the escalation decision, for example when the decision maker was subject to monitoring from another individual. In that condition the responsibility effect was stronger than when one person decided alone, although only four studies examined it [4]. A funding review places a responsible sponsor before a committee that shares authority over the decision. On that limited evidence, adding review to the same people is more likely to harden the case for continuing than to soften it.

A sponsor who asks for one more quarter usually believes in the revised plan. The sponsor has more reason than anyone to believe it, and the committee weighing the request approved the plan now running behind. Giving that committee more authority likely strengthens the people already pulled toward continuing.

What changed when the decision makers left

Students allocating simulated funds face no careers, boards, or auditors, so the harder test is a setting where the money and the reputations are real. Staw, Sigal Barsade, and Kenneth Koput tested the responsibility explanation across 132 California banks over nine years. Turnover among senior bank executives predicted later provisions for loan losses and write-offs of bad loans, while provisions and write-offs did not predict executive turnover [5]. Losses were recognized after senior leadership changed, and recognizing them was not what changed the leadership.

Troubled software projects show a related pattern. Keil and Daniel Robey interviewed forty-two IS auditors about projects that were turned around or sensibly abandoned. In the majority of cases, de-escalation was triggered by actors such as senior managers, internal auditors, or external consultants [6]. The study does not say whether those senior managers had approved the projects; internal auditors and outside consultants ordinarily sit outside a project's approval chain.

Read together, these studies locate the failure. The executive who made the commitment is the person least able to end it, and a gate that asks that executive to argue the case before the committee that shared the approval has built escalation into its own procedure.

Replacing sponsors would be an expensive response. Turnover removes the person who understands the work best, and in the banks, recognition tended to follow a change in leadership. A gate can hand the decision, at a point set in advance, to someone who had no part in the approval.

Designing the stop around the approver

Itamar Simonson and Staw compared de-escalation techniques in a laboratory experiment with 193 business administration students. Three were most effective: making negative outcomes less threatening, setting minimum target levels that lead to a change in policy if they are not achieved, and evaluating decision makers on their decision process rather than the outcome [7]. The second becomes the recorded minimum. The first and third shape how sponsors are judged, a combination the study did not test directly. The field evidence supplies the remaining change: who decides when the program falls short.

EXHIBIT 1

Set the stop at approval, then let the result route the decision

01

At approval

Record the minimum and its date, how the sponsor is judged, and a reviewer with no prior approval.

02

Results arrive

The EPMO checks the reported result against the recorded minimum on the stated date.

03

Minimum cleared

The work continues on its existing terms and the approving committee keeps the decision.

04

Minimum missed

The sponsor recommends continue, narrow, pause, or stop; the named reviewer decides.

Record the minimum result before results arrive

At approval, write down the smallest result that must hold by a stated date for the commitment to continue on its current terms. Use measures the program already reports, such as business units migrated, cost per migration, or unit cost at pilot volume. A miss holds the next round of funding until the continuation decision is made under the rule below.

Simonson and Staw recommend setting such targets before any feedback is received, so later results cannot bias them, and stating them as precisely as possible [7]. Setting the minimum early is what protects it. Once results arrive, everyone who approved the work has a reason to read a shortfall as noise or a delay. Date the first minimum early in the program as well: if the completion effect holds, the pull to continue is weakest while completion is still far off [2].

Keep the minimum separate from the forecast. The forecast is what the sponsor expects. The minimum is the point below which the original case no longer holds, so a program can miss its forecast and still clear its minimum. A missed minimum also works differently from an expiring approval condition, which returns the work to the forum that approved it. At a missed minimum, the decision itself moves.

Judge the sponsor on the quality of the case

Tell sponsors at approval, while the commitment is live, how they will be evaluated: whether the business case was sound given what was known, whether the minimum honestly marked where the case fails, and how early they raised a miss.

In the experiment, telling participants they would be evaluated on the outcome of their original decision did not reduce their commitment, while telling them they would be evaluated on their decision process did [7]. A sponsor credited for raising a miss early can recommend narrowing or stopping without conceding personal failure, which also makes the negative outcome less threatening. When a project stops at its minimum, the record should state whether the original case was sound, and the sponsor's next assignment should follow that judgment.

Let the sponsor recommend and someone else decide

When a minimum result is missed, split the decision. The sponsor prepares the recommendation to continue, narrow, pause, or stop, with the evidence and a revised case. A reviewer, or a small panel, with no prior approval of this commitment or any of its funding rounds makes the decision.

For most programs, a portfolio committee member from outside the sponsoring line of business can serve. For the largest or least reversible commitments, internal audit or an external reviewer fills the kind of role that triggered de-escalation in the projects Keil and Robey studied [6]. The EPMO keeps the minimum on the record, tracks the date, and convenes the reviewer, so the trigger does not depend on the sponsor volunteering the shortfall.

Expect the sponsor's recommendation to lean toward continuing. It comes from a responsible decision maker under review, a condition the meta-analysis associated with a stronger pull to continue, on four studies [4]. The reviewer should weigh it as a well-informed argument from an interested party. The committee that approved the earlier funding rounds stays informed and can ratify the outcome, but it does not make the continue-or-stop call at a missed minimum.

When the work should continue

The reviewer can still decide to continue. Keil and Robey's study covered projects that were successfully turned around as well as those sensibly abandoned [6], and some programs miss an early minimum for a reason that has since been fixed. What the reviewer should decline is an extension on the same case that already failed its own test.

EXHIBIT 2

The evidence locates the failure; it does not certify the design

Method founders
Stage-Gate's founders prescribe tough Go/Kill gates because begun projects are rarely killed.

IS project survey
Some escalation in 30 to 40 percent of projects, with significantly worse perceived outcomes.

Experiments and synthesis
Personal responsibility for the initial decision increases further commitment.

Field studies
Write-offs followed executive turnover; managers, auditors, or consultants triggered de-escalation.

Not established
No study tests this gate design as a whole or shows it works the same way everywhere.

The reviewer knows less than the sponsor, and the design accepts that cost; the sponsor's recommendation carries the detail. For small, reversible commitments the cost is not worth paying, so reserve the split for commitments large enough that a wrong continuation matters.

Sponsors may set minimums they are sure to clear. The approving committee should test whether each minimum marks the point where the case fails, and a minimum that no plausible result could miss should count against the case itself.

Independent review can also turn adversarial, especially if reviewers are rewarded for stops or are inclined to end another leader's program. Evaluate reviewers the way sponsors are evaluated, on the reasoning in the record.

The evidence bounds these claims. The technique comparison comes from one laboratory experiment with business students [7]. The shared-authority finding rests on four studies [4]. The bank study is a field correlation in one industry and one state [5]. Keil, Mann, and Rai relied on auditors' perceptions of project outcomes [2], and Keil and Robey interviewed auditors, who may recall auditor-triggered turnarounds more readily [6]. No study has tested this gate design as a whole.

The operating move

Start at the next funding approval with one commitment large enough to matter. Before the money is released, add to the approval record the minimum result and its date, how the sponsor will be evaluated, and the reviewer who will decide continuation if the minimum is missed.

Then return to the funding review. The sponsor asks for one more quarter to move the largest business unit. The first question is whether the result cleared the minimum recorded at approval. If it did, the work continues on its existing terms. If it did not, the sponsor makes the case and a reviewer who did not approve the funding decides. Nobody deciding has to overturn their own earlier judgment.

Notes and sources

[1] Robert G. Cooper and Scott J. Edgett, Stage-Gate and the Critical Success Factors for New Product Development, BPTrends. https://bptrends.info/wp-content/publicationfiles/07-06-ART-Stage-GateForProductDev-Cooper-Edgett1.pdf

[2] Mark Keil, Joan Mann, and Arun Rai, Why Software Projects Escalate: An Empirical Analysis and Test of Four Theoretical Models, MIS Quarterly 24(4), 2000. https://aisel.aisnet.org/misq/vol24/iss4/4/

[3] Barry M. Staw, Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action, Organizational Behavior and Human Performance 16, 1976, 27-44. https://strategy.sjsu.edu/www.stable/B290/reading/Staw,%20B%20M,%201976,%20Organizational%20Behavior%20and%20Human%20Performance.%2016%20pp%2027-44.pdf

[4] Dustin J. Sleesman, Donald E. Conlon, Gerry McNamara, and Jonathan E. Miles, Cleaning Up the Big Muddy: A Meta-Analytic Review of the Determinants of Escalation of Commitment, Academy of

Management Journal 55(3), 2012, 541-562. http://www.iot.ntnu.no/innovation/norsi-pims-courses/huber/Sleesman,%20Conlon%20&%20McNamara%20(2012).pdf

[5] Barry M. Staw, Sigal G. Barsade, and Kenneth W. Koput, Escalation at the Credit Window: A Longitudinal Study of Bank Executives' Recognition and Write-Off of Problem Loans, Journal of Applied Psychology 82(1), 1997, 130-142. https://experts.arizona.edu/en/publications/escalation-at-the-credit-window-a-longitudinal-study-of-bank-exec/

[6] Mark Keil and Daniel Robey, Turning Around Troubled Software Projects: An Exploratory Study of the Deescalation of Commitment to Failing Courses of Action, Journal of Management Information Systems 15(4), 1999, 63-87. https://www.jmis-web.org/articles/531

[7] Itamar Simonson and Barry M. Staw, Deescalation Strategies: A Comparison of Techniques for Reducing Commitment to Losing Courses of Action, Journal of Applied Psychology 77(4), 1992, 419-426. https://gwern.net/doc/psychology/cognitive-bias/sunk-cost/1992-simonson.pdf

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