

Decision Rights Are a System-Design Choice

Assign who recommends, agrees, decides, and performs before a portfolio — or an agent fleet — starts moving

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

August 2026

EXECUTIVE SUMMARY

Assign the closure rule and follow-through where each consequential decision moves.

A portfolio can have current evidence and still fail to act when nobody knows who may advance, redirect, pause, or stop the commitment. The same open loop appears across an agent fleet when permissions exist without a named closure rule or recovery owner. This paper defines a four-role decision contract and shows how to embed it in the workflow without treating RAPID or RACI as doctrine.

- 01** Write consequential decisions in verb form, including the choice, boundary, trigger, and operating consequence.
- 02** Separate Recommend, Agree, Decide, and Perform; keep veto rights narrow and require an alternative or escalation path.
- 03** Name one accountable closure rule while preserving collective authority, separated approvals, appeal, and contestability where required.
- 04** For agent fleets, keep rights and recovery accountability with named people or governance bodies; delegate preparation, enforcement, and execution within explicit local and aggregate limits.

CONTENTS

The open loop	The agent-fleet application: authority must travel with the action
The decision is the unit of design	A single decider is not a complete system
A four-role decision contract	Match the weight of governance to the door
Choose the decider by consequence, context, and conflict	Measure closure, not the number of matrices
Put the contract where the work moves	The design review
The portfolio application: rights across the system	Close the loop on purpose
	Notes and sources

The open loop

A portfolio review reaches the hard question. The evidence says an initiative is late, its assumptions have weakened, and another priority now needs the same people. Everyone can see the conflict. No one in the room is sure who can redirect

the work, pause the next release, or stop the commitment. The meeting ends with a request for more analysis.

Nothing is missing from the dashboard. What is missing is authority attached to the decision. The evidence has reached the organization, but the control loop remains open because no one has been assigned the right and obligation to act on it.

The defect sits in the operating design. The workflow produces information and activity without converting a consequential choice into a named act. Titles and committees exist, and accountability appears on an org chart, yet the decision itself has no closure rule.

Gary Neilson, Karla Martin, and Elizabeth Powers studied a database of more than 125,000 employee profiles across more than 1,000 organizations in over 50 countries. Decision rights and information flow were the two strongest execution levers in their analysis.[2] A freely readable account by former Harvard Business Review editor Thomas A. Stewart reports the same sample and finding.[9] The pairing matters. Authority without information is arbitrary. Information without authority is inert.

The operating standard is therefore not simply, "make accountability clear." It is: for each consequential recurring decision, assign who develops the recommendation, who has a narrow right to agree or veto, who makes the decision, and who performs what follows. Then place those rights inside the gate, system, and record where the work moves.

The decision is the unit of design

Organizations usually assign authority by role. A business-unit leader owns a function. A sponsor owns a program. A steering committee oversees a portfolio. Those statements may be true and still fail to answer the question in front of the team: who can approve the next funding tranche, resolve a risk exception, accept a release, redirect capacity, or stop an agent from acting?

An org chart describes durable reporting relationships. A consequential decision crosses those relationships. It draws on evidence from delivery, finance, operations, risk, architecture, legal, and the people who must execute the result. If the organization relies on the chart alone, each participant reads authority differently. The sponsor assumes the committee will decide. The committee assumes the executive sponsor retains the call. Control functions assume they can block the action, but no one has defined what a block requires or where disagreement goes.

The fix begins by naming the decision in verb form. "Portfolio governance" is not a decision. "Release the next six months of capacity" is. "AI oversight" is not a decision. "Allow this agent to issue a customer credit up to a stated limit" is. A useful decision statement contains the choice, the boundary, the trigger, and the consequence.

- Choice: advance, redirect, pause, stop, approve, reject, or escalate.

- **Boundary:** the funding, capacity, customer impact, legal exposure, or agent authority covered.
- **Trigger:** the evidence, date, exception, or threshold that puts the choice on the table.
- **Consequence:** what changes in the workflow when the decision is recorded.

That formulation turns a vague responsibility into an actuator. It also prevents a common mistake: assigning one person to "own" a domain while leaving every material cross-functional decision to ad hoc negotiation.

A four-role decision contract

RAPID, RACI, DACI, and similar models are useful because they force distinctions that ordinary responsibility language blurs. They are not doctrine. The label matters less than whether the workflow separates contribution, consent, choice, and execution. Paul Rogers and Marcia Blenko's RAPID model names five roles: Recommend, Agree, Perform, Input, and Decide. Bain also warns that the framework is not a panacea; a weak or indecisive decision-maker can defeat a well-drawn map.[1]

EXHIBIT 1

Four rights carry a choice from evidence to changed work

Recommend

Build the decision-complete case, alternatives, evidence, limits, and deadline.

Agree

Exercise a bounded veto for a named obligation, with an alternative or escalation.

Decide

Apply the accountable closure rule and record the choice, conditions, and rationale.

Perform

Change the work, funding, permission, or record, then report any exception.

A compact four-role contract fits recurring portfolio and agent-fleet choices when one recommender can gather the relevant input, any veto can be bounded, and the decision has a clear closure rule. Input is organized through the recommendation instead of becoming a parallel claim to the decision. Keep Input explicit where law, collective governance, statutory consultation, separation of duties, or an independent challenge function requires its own visible standing. The model should expose those conditions rather than flatten them.

Recommend

The recommender owns the decision case. This role frames the choice, gathers the required evidence, shows alternatives, exposes uncertainty, and states a recommendation. The recommender is not a neutral meeting coordinator. The role has a deadline and a burden of proof.

A recommendation should be decision-complete: the decision requested, options considered, evidence used, consequences, known limits, and the latest useful date. If the recommender cannot produce that package, the decision is not ready. If the recommender produces it and the meeting still cannot identify a decider, the governance design is not ready.

Agree

Agreement is a narrow control right, not a courtesy copy and not a license for every affected function to demand consensus. An agree role exists where a function carries a non-delegable obligation or a material exposure: law, safety, regulatory compliance, a binding architecture boundary, or another clearly stated constraint.

Veto rights must be explicit. They also need a rule of use. Bain's account says only a few roles should hold veto power and that a veto should come with an alternative or an escalation to the decider.[1] Without that discipline, "agree" becomes an unbounded queue of silent stops. The organization gains the appearance of control while accountability diffuses across everyone who can withhold consent.

Decide

The decision needs one accountable closure rule. In ordinary operating choices, that rule should name one person. Law, fiduciary duty, labor agreements, safety authority, or governance charters may instead require a board or another collective body. The contract then names how that body closes the vote, records the outcome, resolves a tie, and handles a missed deadline. "The committee decides" remains incomplete if the committee can adjourn without a disposition.

A single decider does not mean unilateral judgment. The decider must use the recommendation, hear required agreement roles, weigh material input, and accept the consequence of the choice. The design goal is a single point of closure, not a heroic executive.

Perform

The performer translates the decision into changed work. Funding moves. A backlog changes. A release is held. An agent's permission is revoked. A customer case is recovered. The performer confirms completion and reports exceptions back to the decision record.

This role is what keeps governance from ending at approval. A decision without a named performer is an opinion with minutes. The work continues under its previous state until someone changes the system of record, the resource commitment, or the operating permission.

Choose the decider by consequence, context, and conflict

Naming one decider forces a second question: which person should hold the right? Seniority is a weak default. The best position depends on where the consequence lands, where the relevant context sits, and who can resolve the trade-off without becoming captured by one side of it.

Start with consequence and context. A capacity decision that displaces business commitments belongs with the leader accountable for those operating outcomes, not with the analyst who assembled the portfolio view. That leader still needs local information before the latest useful date. When the decider sits far from the work, strengthen the evidence package and the route for clarifying questions rather than pushing the choice downward by default.

Then test conflict. A sponsor defending the work may be well placed to recommend continuation but poorly placed to decide whether another investment should take its capacity. A control function can veto a prohibited action without choosing the business alternative. Separating those roles makes the tension inspectable.

Finally, test capacity. Hundreds of nominal rights turn an executive into a queue. Narrow the decisions that rise to that role, delegate within stated thresholds, or give a standing forum a reliable closure rule. The chosen arrangement must operate at the cadence the work requires.

These tests do not produce a universal answer. They produce an inspectable assignment. The organization can explain why this decider holds this choice, what evidence must reach that person, which constraints can block it, and when the right should move.

Put the contract where the work moves

A role map in a presentation is documentation. A role map inside the workflow is governance. The difference is whether the work can cross a consequential boundary without the required decision.

EXHIBIT 2

The workflow must carry the decision through to recovery**01****Trigger**

Evidence, a threshold, a date, or an exception puts the choice on the table.

02**Recommendation**

The named owner presents a decision-complete case before the latest useful date.

03**Bounded agreement**

Required functions exercise only the veto attached to their stated obligation.

04**Closure**

The named person or collective rule records advance, redirect, pause, stop, or escalate.

05**Perform or recover**

The owner changes the system and routes missed conditions or failure into recovery.

At portfolio intake, the record should name the decision being requested, the recommender, required agree roles, the decider, and the performer before the item reaches a review. At a stage gate, the agenda should show the permitted dispositions and the person who can choose among them. In delivery tools, the decision record should carry the evidence date, decision deadline, outcome, conditions, and next review. In an agent platform, the permission boundary should reflect which actions are autonomous, which require sign-off, which trigger escalation, and which can be stopped immediately.

The workflow also needs a disagreement path. If an agree role rejects the recommendation, the record should capture the specific constraint, the acceptable alternative, and the escalation destination. If the decider misses the latest useful date, the system should not silently continue. It should apply a pre-agreed default: pause the commitment, escalate the choice, or allow a time-bounded continuation with visible risk acceptance.

These mechanics do more than clarify responsibility. They preserve the decision's state. People joining later can see what was authorized, on what evidence, under which conditions, and who must revisit it. That is essential when a portfolio spans quarters or an agent acts thousands of times between governance meetings.

Project Management Institute's 2026 systems research describes contemporary project work as more interconnected and less predictable, with complexity appearing as decision friction and rework.[4]

That is the environment in which informal authority fails most often. More stakeholders and faster signals do not close the loop. They create more paths through which the call can diffuse.

The portfolio application: rights across the system

A gate answers whether one commitment may cross one boundary. A decision-rights system governs the recurring choices that surround the whole portfolio. The distinction matters. The authority to stop at a gate is one decision. The authority to admit demand, sequence capacity, accept an exception, redirect funding, or recognize a benefit sits elsewhere in the same system.

Stage-Gate's portfolio work describes portfolio management as a dynamic decision process in which projects are evaluated, selected, prioritized, accelerated, killed, or deprioritized and resources are reallocated.[3] Each verb is a decision. Treating "the portfolio committee" as the owner of all of them creates a meeting-heavy center that either becomes a bottleneck or pushes real choices into pre-meetings.

A better design maps rights per recurring choice. A portfolio office may recommend sequencing because it sees demand and capacity across the system. A business executive may decide because the choice changes operating outcomes. Risk or architecture may hold an agree right for a bounded exposure. Delivery leaders perform the redirect by changing staffing, scope, milestones, and systems of record.

The roles can change by decision. The person who decides whether to admit an initiative may not be the person who accepts a production risk. The sponsor who recommends continuation should not automatically control every later decision about further commitment. The point is not to centralize power in the PMO. It is to make the route from evidence to authorized action visible.

The contract also routes information. Current capacity, risk, outcome, and dependency evidence must reach the recommender; agree roles receive the evidence relevant to their obligation; the decider sees the trade-off before the latest useful date; and the performer receives a disposition precise enough to execute. Stale evidence or a missing decider breaks the same chain at different points.

Operating illustration: a capacity conflict

Consider a portfolio with two active initiatives that need the same specialist team. One initiative is behind plan but remains strategically important. The other has reached a regulatory deadline. A generic steering committee can discuss the collision for weeks because each sponsor owns only one side of it.

A decision contract changes the path. The portfolio office recommends a sequence using dependency, deadline, capacity, and consequence evidence. The regulatory owner holds an agree right over any option that misses the non-delegable obligation. The operating executive who owns the combined

outcome decides the allocation. Delivery leaders perform the decision by changing team assignments, release scope, and dates in the systems of record.

The contract does not guarantee a painless result. One initiative still loses time. Its value is that the trade-off is made at the level that owns the combined consequence, with the constraint visible and the follow-through assigned. Delay can no longer masquerade as neutrality.

The agent-fleet application: authority must travel with the action

Agent governance often starts with a human-in-the-loop statement. That is not a decision-rights design. It does not say which human, for which action, at what threshold, with what evidence, or what happens when the human does not respond.

McKinsey's agentic-organization work argues that governance for continuously operating agents must become embedded and data driven while humans retain final accountability.[5] Deloitte makes the decision boundary more concrete: leaders should define what agents may decide, what requires human sign-off, what must escalate, what evidence an agent must produce, and who resolves exceptions. It also calls for stop mechanisms where errors could spread before people notice.[6]

The four rights remain with named people or organizational bodies. Agents and rules can prepare, enforce, or execute a delegated policy, but software does not acquire organizational accountability from a permission setting.

- **Recommend:** a named service or product owner holds the recommendation right; an agent may assemble the evidence and propose an action with sources, confidence limits, and policy checks attached.
- **Agree:** a named risk, legal, safety, or policy function holds a bounded veto. A rule or control agent may enforce its pre-authorized constraint and route the exception to that function.
- **Decide:** a named human or governance body sets the authority band and closes consequential exceptions. Automation applies the already-authorized policy inside that band.
- **Perform:** a named operations or service owner remains accountable for follow-through and recovery. An agent may execute the action and record what changed.

The crucial design move is to attach authority to the action, not to the agent as a whole. One agent may recommend a refund, perform a refund under a low threshold, and have no right to decide a larger credit. Another may stop a workflow when a safety rule fires but lack the authority to restart it. Agent identity alone is too coarse; decision type, consequence, and operating state determine the right.

Static responsibility charts are also insufficient when agents learn, tools change, or permissions expand. Deloitte's 2026 human-capital work says AI decision rights need override privileges, escalation paths, and consensus rules engineered into the system.[7] McKinsey and Deloitte provide

practitioner guidance for designing those boundaries; they do not validate this paper's four-role adaptation. The map therefore remains a testable operating choice. Changes to models, tools, thresholds, or data access trigger a review of the decisions the fleet is authorized to make.

Fleet composition adds another risk. Three actions may each sit inside an agent's local authorization band yet combine into a customer, financial, or operational outcome that no one intended to authorize. Aggregate thresholds must therefore span the fleet: cumulative credits, repeated retries, linked data changes, or correlated exceptions can trigger a shared stop. One named owner decides whether the fleet restarts and remains accountable for recovery across the chain.

This paper therefore differs from a supervision model for one agent. The unit here is the network of consequential decisions across the fleet: who may propose, block, authorize, execute, and recover each class of action. A named human remains accountable where the organization must answer for the consequence, but not every action requires the same human review.

Operating illustration: a customer-action fleet

Imagine a service workflow in which one agent classifies a complaint, another recommends a remedy, and a third applies the approved action. The system is fast, but speed alone says nothing about authority. If all three agents inherit broad permissions from the same service account, the organization cannot tell which decision was delegated and which was merely technically possible.

A decision-rights design separates the action classes. The classification agent supplies input. The recommendation agent may propose a credit and attach the policy basis. A policy control enforces a pre-authorized stop outside contract terms and routes the case to the named policy owner. The fulfillment agent may perform credits below a tested threshold after the recommendation passes checks. A named service leader decides exceptions above that threshold, and a recovery owner can reverse an incorrect credit, contact the customer, and adjust the rule.

Logs then capture more than agent activity. They capture the decision chain: proposal, evidence, veto or agreement, authorization, performance, and any recovery. That record lets a reviewer determine whether the system acted within its rights rather than merely whether the software ran as designed.

A single decider is not a complete system

Clear authority can still produce bad outcomes. The decider may lack competence. Incentives may reward continuation. Evidence may arrive late. Agree roles may protect narrow interests. Performers may be unable to execute. A role map can even make a broken design look disciplined because every box has a name.

EXHIBIT 3

Clear rights close one part of the operating loop**Supported**

Bain distinguishes decision roles; Neilson, Martin, and Powers link rights and information flow.

**Portfolio use**

Stage-Gate treats selection, priority, stopping, and resource movement as recurring decisions.

**Agent guidance**

McKinsey and Deloitte call for embedded boundaries, escalation, overrides, and human accountability.

**Required beside it**

Current evidence, capable owners, legitimate challenge, executable work, and recovery paths.

**Not established**

No universal single decider, causal performance claim, or validation of the four-role model.

Bain's warning that RAPID is not a panacea should be taken literally.[1] Decision rights close one part of the loop. The rest still needs a trustworthy signal, a capable owner, an executable recovery path, and a commitment that can be changed. Weak evidence calls for a conditional decision. Irreversible consequences deserve more review, and approval should wait when the performer lacks capacity.

Collective authority can be the correct design. Boards, fiduciary committees, safety bodies, works councils, and separated approval functions carry mandates that should not collapse into one executive. Appeal and contestability may also require a second, independent decision path. The contract still needs closure, but it preserves plural authority by naming the body, quorum, decision rule, challenge right, escalation, and record rather than forcing consensus into a fictional individual.

Overload is a separate design risk. If every choice rises to the same executive, formal clarity creates a bottleneck. A decision inventory should set thresholds for choices that remain local, cross a gate, operate through delegated guardrails, or require escalation.

This is why the system must distinguish authority from advice. Broad input can improve the recommendation. Required agreement can protect a genuine constraint. Neither should create an

informal veto. The decider must be able to close the choice, and the record must show whose advice was not followed and why when the disagreement is material.

Match the weight of governance to the door

Formal decision rights earn their cost when consequence and irreversibility are high. They become bureaucracy when copied onto routine choices that are cheap to undo. Amazon's 2015 shareholder letter describes this as the difference between one-way-door decisions, which deserve deliberation, and two-way-door decisions, which should move quickly through high-judgment individuals or small groups.[8]

Test consequence and reversibility instead of seniority. A low-value agent action inside a tested limit may execute automatically, while a change to that authority boundary returns to a named human. Delivery owners can keep small project adjustments local. A capacity shift that displaces three other commitments belongs at the portfolio level because the consequence crosses those commitments.

Proportionality keeps the model usable. Apply the short contract where ambiguous authority creates material delay, hidden risk, or unowned action; leave routine, reversible work local.

Measure closure, not the number of matrices

A decision-rights rollout can become another documentation program if success is measured by completed role charts. The operating measures should test whether consequential choices now close and change work.

Start with decision latency: elapsed time from a decision-complete recommendation to a recorded disposition. Separate waiting for missing evidence from waiting for authority. That distinction tells leaders whether the repair belongs in the recommendation standard or the decision path.

Track silent defaults. Count cases in which work continued because no decision arrived by the latest useful date. A continuation may be the right answer, but it should be an active, recorded choice or a pre-agreed time-bounded default. Otherwise delay becomes an invisible authorization.

Track veto behavior. A healthy agree role uses its authority rarely, cites the violated constraint, and supplies an alternative or escalation. Repeated vetoes without alternatives point to an unbounded control right. Repeated escalations point to a threshold or capability problem that the role map has exposed but not solved.

Track execution and recovery. Did the performer change the system within the agreed time? Did the decision produce an exception or reversal? For agent fleets, did actions stay inside their authorization

bands, and did overrides reach the named decider? These measures do not prove the decision was good. They show whether the designed path operated and where it broke.

Finally, review the decision inventory itself. New products, regulations, tools, agents, and operating dependencies create new choices. Other decisions become routine enough to delegate. The map should change as consequence and reversibility change. A static decision-rights catalog will drift back into the same ambiguity it was built to remove.

The design review

A portfolio or agent-fleet leader can test the operating system by sampling five consequential decisions from the last quarter. Do not start with the org chart. Start with the actual choice and trace what happened.

- Could the team state the decision in one sentence, including its boundary and consequence?
- Was one recommender responsible for a decision-complete case by a useful date?
- Were agree and veto rights narrow, explicit, and tied to a real obligation?
- Could one named decider close the choice without waiting for informal consensus?
- Did a named performer change the work, permission, funding, or record after the call?
- If the decision was late or wrong, did the workflow trigger escalation and recovery rather than silent continuation?

A "no" exposes a design gap. Several no answers around the same function expose a systemic one. The repair is specific: change the role, evidence requirement, threshold, escalation, or workflow state. Another accountability speech will not do it.

Close the loop on purpose

Governance becomes real when evidence can change action. The organization may know that a project should stop, capacity should move, or an agent crossed a boundary. Without a closure rule and assigned follow-through, the system has observed the problem without controlling it.

Design decision rights per consequential choice, carry them in the workflow, and revise them when the operating context changes. Recommend, agree, decide, and perform are a practical test rather than doctrine. Apply that test to five recent decisions and repair the role, threshold, evidence requirement, challenge path, or workflow state wherever the chain broke.

Notes and sources

- [1] Paul Rogers and Marcia Blenko, "Who Has the D?" Bain & Company, June 28, 2006. <https://www.bain.com/insights/who-has-the-d/>
- [2] Gary L. Neilson, Karla L. Martin, and Elizabeth Powers, "The Secrets to Successful Strategy Execution," Harvard Business Review 86, no. 6 (June 2008): 60-70. No reader-facing link is included because the full article is subscription-only.
- [3] Robert G. Cooper and Scott J. Edgett, "Portfolio Management: Fundamental for New Product Success," Stage-Gate International. <https://www.stage-gate.com/blog/portfolio-management-fundamental-for-new-product-success/>
- [4] Project Management Institute, Pulse of the Profession 2026: Driving Success in Complex Projects - From Navigating Tasks to Navigating Systems, May 2026. <https://www.pmi.org/learning/thought-leadership/driving-success-in-complex-projects>
- [5] Alexander Sukharevsky, Alexis Krivkovich, Arne Gast, Arsen Storozhev, Dana Maor, Deepak Mahadevan, Lari Hämäläinen, and Sandra Durth, "The Agentic Organization: Contours of the Next Paradigm for the AI Era," McKinsey & Company, September 26, 2025. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>
- [6] David Mallon, Brad Kreit, and Natasha Buckley, "Rethinking Operating Models for Humans with Agents," Deloitte Insights, April 2, 2026. <https://www.deloitte.com/us/en/insights/topics/talent/operating-models-for-humans-ai-agents.html>
- [7] David Mallon, Julie Duda, Stefano Besana, and Maya Bodan, "AI and the Future of Human Decision-Making," Deloitte Insights, March 3, 2026. <https://www.deloitte.com/us/en/insights/topics/talent/human-capital-trends/2026/decision-making-with-ai.html>
- [8] Jeffrey P. Bezos, Amazon 2015 Letter to Shareholders, 2016. https://s2.q4cdn.com/299287126/files/doc_financials/annual/2015-Letter-to-Shareholders.PDF
- [9] Thomas A. Stewart, "Putting Strategy into Practice," strategy+business, April 19, 2010. <https://www.strategy-business.com/article/00030>

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

This white paper may be shared freely with attribution. © 2026 Marco Policani.