WHITE PAPER · PORTFOLIO & AI GOVERNANCE

AI Adoption Starts With the Constraint, Not the Tool

A portfolio-governance standard for turning AI enthusiasm into better recurring work

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

July 2026

Executive premise

AI adoption should begin with the constraint that limits value, not with the tool that is easiest to demonstrate.

The distinction sounds small until money, scarce attention, and change capacity are involved. Tool-first adoption asks what a model, platform, agent, or vendor product can do. Constraint-first adoption asks where recurring work is blocked, what outcome would improve if the blockage moved, and which intervention has the best chance of changing that condition without adding unnecessary complexity.

That intervention may be AI. It may also be a new workflow, better source data, clearer decision rights, training, deterministic automation, a vendor feature already available in the stack, a governance control, a new role, or a decision to wait. The practical job of AI portfolio governance is to make those options comparable before the organization commits funding and credibility to one path.

The market signal supports this discipline. McKinsey reported in 2025 that 88 percent of respondents were using AI in at least one business function, while only 39 percent reported enterprise-level EBIT impact. The same survey found broad experimentation with AI agents and a strong link between high performance and workflow redesign.1 The gap is widest around agents: McKinsey’s own follow-up found that AI is everywhere but the agentic organization is not yet, with adoption running ahead of the operating redesign that makes agents dependable.4 The pattern is not lack of interest. It is a conversion problem: activity becomes value only when it attaches to work that leadership actually needs to change.

RAND reached the same problem from the builder side. Its 2024 analysis of AI project failures found that stakeholders often misunderstand the problem to be solved, optimize for the wrong metric, or chase the latest technology instead of a real user problem.2 Gartner updated the operational warning in 2026, finding that at least half of GenAI projects had been abandoned after proof of concept by the end of the previous year because of poor data quality, inadequate risk controls, escalating costs, or unclear business value.3

Those findings do not mean organizations should slow AI adoption to a crawl. They mean adoption needs a better first gate. A tool is not a strategy. A pilot is not proof. A demo is not a use case. The first governable unit is the constraint.

Why tool-first adoption creates noise

Tool-first adoption has an appealing rhythm. A vendor or internal team shows what the technology can do. Leaders see speed. Teams identify tasks that feel similar to the demo. Licenses spread. Usage dashboards appear. The organization can point to activity within a quarter.

The trouble is that activity is rarely the thing the portfolio funded. Leadership funds AI to reduce rework, accelerate cycles, increase analytical coverage, improve decisions, lower cost to serve, reduce risk, or open a new growth path. Those outcomes live in workflows, not in the tool console.

A tool-first pattern also hides three separate questions inside one conversation. Is the technology capable? Is the workflow ready? Is the intervention worth the operating change? A capable model can still fail in a poorly governed workflow. A promising workflow can still fail without usable data or decision rights. A valuable idea can still lose to a cheaper non-AI intervention.

The effect is portfolio noise. Teams compete for scarce AI attention with requests that cannot be compared: a chatbot for one group, a summarizer for another, an agent proof of concept, a data-science idea, a training need, a compliance concern, a vendor upgrade. Without a constraint frame, the loudest or most fashionable request often wins the next increment of capacity.

Constraint-first governance gives the portfolio a common unit of comparison. It asks every idea to name the bottleneck, the affected recurring workflow, the measurable condition that would improve, the owner of the value claim, the evidence required, and the smallest intervention worth trying. That does not remove judgment. It makes judgment inspectable.

The constraint-first standard

A constraint is the current condition that limits value in recurring work. It can be a slow handoff, missing data, inconsistent classification, unclear approval rights, poor exception handling, manual synthesis, customer friction, policy ambiguity, review capacity, or a lack of trust in available evidence.

The standard has four parts. First, name the recurring workflow and the specific constraint. Second, describe the value mechanism: what changes if the constraint moves. Third, compare intervention types before picking technology. Fourth, define the evidence that would justify more reliance, more funding, or a broader rollout.

The discipline is deliberately plain. It prevents teams from treating AI as a default answer while still giving AI a serious path to funding when it is the right answer. The point is to stop treating novelty as readiness.

A constraint-first intake should ask a short set of questions. What work repeats often enough for improvement to matter? Where does the work slow down, break, rework, or require scarce judgment? What baseline exists now? Which human role owns the outcome? What would count as credible improvement? What could go wrong if the intervention is trusted too early? What non-AI option should be compared?

These questions fit naturally with NIST AI RMF language. Govern, map, measure, and manage create a cycle for understanding context, assessing performance and risk, and deciding how the system should be used over time.5 Constraint-first adoption turns that cycle into a portfolio habit.

The minimum effective intervention

The minimum effective intervention is the least complicated change that can credibly move the constraint at the needed level of reliability, speed, risk, and cost.

This is where many AI conversations become clearer. If the problem is inconsistent source data, a model may only make the inconsistency faster. If the problem is unclear decision rights, an agent can route work without resolving who can accept risk. If the problem is review capacity, AI may help prepare evidence, but the governance question is still who is allowed to rely on it.

Minimum does not mean small or cheap by default. Some constraints deserve serious investment because they sit on a critical path or reshape a profit pool. McKinsey argued in 2026 that AI value will not be created evenly; productivity gains become table stakes, while deeper value comes from reshaping offerings, business models, and market structures.7 A high-value constraint may justify major redesign. The discipline is to make that call explicitly.

The minimum effective intervention compares at least six response types: redesign the workflow, clean or govern the data, train or reassign people, automate deterministic steps, buy or configure a tool, or use AI to augment or automate part of the work. A strong portfolio does not reward the most AI-shaped answer. It rewards the answer that moves the constraint with acceptable evidence and operating cost.

This is the practical bridge between innovation and governance. Teams can experiment quickly, but every experiment has a value hypothesis, an evidence path, and a decision point. If the evidence is weak, the idea is refined or stopped before it becomes a sunk-cost story.

Sequence matters as much as selection. Data readiness precedes model reliability, decision rights precede automation, and process clarity precedes augmentation. An intervention placed ahead of its prerequisite fails in a way that looks like a technology problem but is not. A classifier trained on inconsistent source data produces confident inconsistency. An agent handed a workflow with no accountable owner routes work faster than anyone can answer for it. The minimum effective intervention is therefore not only the smallest change that moves the constraint, but the one whose prerequisites already hold — or the prerequisite itself, sequenced first.

Reading the constraint

Most constraints announce themselves as a symptom in the work, not as a technology gap. The discipline is to read the symptom back to its cause before choosing a response, because the same symptom can hide different constraints and each one routes to a different intervention. The table below pairs common symptoms with the constraint they usually signal, the response that tends to move it, and the AI reflex that looks productive and changes nothing.

Symptom in the work	What the constraint usually is	Intervention that moves it	The wrong AI reflex
Status takes days to assemble and is stale on arrival	Manual synthesis across disconnected systems	Restructure intake, then use AI to prepare exceptions for human review	Point a summarizer at the same messy inputs
The same request is classified three different ways	Absent or unenforced decision rules	Define the rules, then automate the deterministic part	Ask a model to reinvent the taxonomy each time
Approvals stall for weeks	Unclear decision rights, not slow drafting	Assign the decision right and the escalation path	Add an agent that routes work no one is authorized to accept
Analysts cover only a fraction of the cases	Review capacity	Use AI to widen coverage while humans keep the judgment	Treat model output as the decision and drop the review
The forecast is never trusted	Poor or ungoverned source data	Fix data ownership and quality first	Layer a model on data no one believes

The right column is consistent. AI applied to an unresolved data, decision-right, or process constraint accelerates the symptom instead of removing it. The response that moves the constraint is often not the AI-shaped one, and a portfolio that cannot tell the difference will fund motion in place of progress.

A portfolio operating model

Constraint-first adoption needs a light operating model. The model can run through five gates.

Gate one is constraint intake. The idea owner names the workflow, the current bottleneck, the baseline, the user group, and the decision or outcome that should improve. Ideas without a named workflow stay in exploration, not funding.

Gate two is intervention routing. A small review group compares likely responses: workflow, data, people, governance, deterministic automation, vendor feature, model, agent, or wait. The output is not a full business case. It is a routing decision with a reason.

Gate three is proof design. The team defines the smallest credible test, the technical and operating guardrails, the adoption signal, and the workflow measure that would matter. McKinsey's 2026 measurement framework is useful here because it separates technical performance, user adoption, operational KPIs, strategic outcomes, and financial impact.6 The early gate should not pretend to prove financial impact before the operational signal exists.

Gate four is reliance decision. The organization reviews the evidence and decides whether to scale, refine, stop, or switch interventions. This is where many pilots fail quietly: the demo worked, but adoption did not enter the critical path, the business measure did not move, or the total cost of ownership became unattractive.

Gate five is value realization. If the intervention scales, it enters normal performance management. Owners continue to review adoption, quality, cost, exception rates, and drift. AI work is not done when the tool ships. It is done when the changed workflow keeps producing the value case under real conditions.

What it looks like in practice

The safest public examples are operating patterns, not claims of broad enterprise AI ownership. In portfolio governance work, the constraint often appears as decision friction: too many parallel facts, too many dependencies, and too little time in the cadence for the decision that matters.

In one published case pattern, the useful AI role was to prepare the evidence so the human cadence could spend less time reconstructing status and more time deciding what needed attention. The bottleneck was decision-quality time, not document production. AI helped decompose plan data, surface inconsistencies, and focus human review on critical-path issues.

In another pattern, the constraint was classification quality across a large initiative set. AI could propose reclassifications, but product managers still confirmed, declined, or amended the result. The value came from attaching AI to a defined recurring workflow with named human ownership and an evidence trail.

These examples illustrate the standard. Start with the bottleneck. Use AI where it changes the work. Keep humans accountable for the decision. Measure the workflow outcome rather than the excitement around the tool.

A worked example

The following is an illustration, not a client account. It shows how one constraint moves through the five gates.

A shared-services team spends the first half of every month assembling a status view across a dozen delivery streams. Leaders receive it late, and by the time they read it, the picture has already moved. The tool-first request arrives on schedule: license a generative assistant to write the report faster.

At constraint intake, the owner names the workflow and the bottleneck. The workflow is the monthly portfolio review. The constraint is not report production; it is the time leaders spend reconstructing what happened instead of deciding what to do. The baseline is measurable: reviews open with roughly forty minutes of status recitation before the first real decision.

At intervention routing, the group compares responses. A summarizer would produce a faster version of the same recitation. The higher-value move is to restructure intake so status arrives in defined fields, then use AI to flag only the inconsistencies and critical-path risks that need attention. The routing decision is explicit: redesign the intake, apply AI to exception preparation, defer any autonomous agent.

At proof design, the team sets a small test: two review cycles, a fixed set of streams, one operating measure (decision time in the meeting) and one quality check (whether the flagged exceptions were the ones that actually mattered). The reliance boundary is stated up front. AI prepares the evidence; the portfolio lead owns every call.

At the reliance decision, the evidence is read. Decision time moved, and the flagged exceptions matched the issues leaders would have raised on their own. The team scales to the full stream set and keeps the named human owner in place.

At value realization, the workflow enters normal management. The measure is not how often the assistant runs. It is whether the review still opens with a decision instead of a recitation. Nothing in this path depended on the specific model. It depended on naming the constraint, choosing the smallest intervention that moved it, and keeping a named human accountable for the result.

How a steering committee can use the standard

The standard works best when it is used before a team has fallen in love with a solution. A steering committee does not need to become an AI architecture board. It needs a repeatable way to turn AI enthusiasm into comparable operating choices.

A simple review sequence is enough. The sponsor names the constraint and the recurring workflow. The process owner confirms the current baseline and explains how the work is governed today. Technology or data leads describe the available intervention options and the readiness of the data. Risk, legal, security, or compliance partners identify reliance boundaries. Finance or portfolio operations asks what evidence would justify the next funding increment.

That conversation changes the shape of the meeting. The committee no longer has to approve or reject AI in the abstract. It can decide whether a specific constraint is worth moving, whether the proposed intervention is proportionate, and whether the proof plan is strong enough for the next stage.

This also gives leaders a cleaner way to stop work. A stopped pilot is not automatically a failure. It may be evidence that the original constraint was poorly framed, the data was not ready, the workflow owner was missing, or a simpler intervention should be tried first. Governance creates that off-ramp before a weak idea becomes politically expensive.

A practical scorecard

A constraint-first scorecard should be short enough to use in a real portfolio review. Five questions usually cover the decision.

First, is the constraint specific? A vague problem such as productivity or knowledge work will produce vague AI demand. A useful constraint names where the work slows, breaks, repeats, or consumes scarce judgment.

Second, is the value mechanism credible? The team should be able to explain how the intervention would reduce cycle time, lower rework, improve quality, broaden analytical coverage, reduce risk, increase capacity, or change a customer or business outcome.

Third, is AI the minimum effective intervention? The answer may be yes. But the team should have compared workflow redesign, data cleanup, training, deterministic automation, vendor features, and a governed wait.

Fourth, is the proof path clear? A team should know the baseline, the adoption signal, the workflow metric, the human review point, and the decision gate before it scales beyond a small test.

Fifth, is ownership named? A tool owner is not enough. The work needs a process owner for the outcome, a technical owner for the system, and a governance owner for the reliance boundary.

Scoring these questions is less important than forcing the conversation. The value is in making the assumptions visible while the choice is still reversible.

Common failure patterns

The first failure pattern is the demo-to-roadmap jump. A team sees a plausible demonstration and immediately asks for rollout funding. The missing middle is proof that the same result holds in the organization's workflow, data, controls, and adoption reality.

The second pattern is the orphaned pilot. A team runs an experiment without a process owner who can change the work if the experiment succeeds. The pilot can produce interesting output, but it has no path into business as usual.

The third pattern is the telemetry trap. Usage rises, but the workflow does not improve. This is especially tempting when platform dashboards are easier to obtain than operational measures. The portfolio should treat usage as a supporting signal, never as the value claim by itself.

The fourth pattern is overfitting to the hardest model. Frontier models are useful for exploration, synthesis, and uncertain tasks, but many scaled workflows need reliability, cost control, and repeatability more than maximum model capability. Model choice is part of the intervention decision.

The fifth pattern is unmanaged reliance. The system starts as an assistant, then its outputs quietly become decision inputs without new review rights, quality sampling, escalation rules, or auditability. Constraint-first governance should name the intended reliance level before the team confuses convenience with authority.

The evidence pack

A constraint-first evidence pack should be compact enough to survive real executive review. It does not need to become a research binder. It needs to let leaders see the connection from problem to intervention to proof.

The first page names the constraint, the affected workflow, the current baseline, and the value mechanism. It should be written in operating language, not vendor language. A reader should understand what work will change even if the AI tool name is removed.

The second page compares interventions. It should show why the team chose AI, a workflow redesign, a vendor feature, a deterministic automation, a data cleanup effort, training, or waiting. This page is where many weak projects become visible. If the team cannot explain why the chosen answer is better than the simpler alternatives, the portfolio should slow down.

The third page defines the proof plan. It names the pilot population, the technical signal, the adoption signal, the workflow measure, the quality review method, the reliance boundary, and the decision gate. This keeps pilots from drifting into open-ended experimentation.

The fourth page records the decision. Scale, refine, stop, switch intervention, or wait. The decision should include a date, owner, next evidence point, and explicit reason. That record becomes valuable later when leaders ask why the portfolio funded one AI idea and not another.

Decision rights and human accountability

Constraint-first governance also clarifies who is allowed to accept the result of AI-supported work. Many organizations say they want human review, but review without decision rights is weak protection.

The process owner should own the business outcome. The technical owner should own system health, integration, and monitoring. The governance owner should own the reliance rule: where AI can draft, recommend, classify, summarize, execute, or escalate, and where a named human must decide.

Those roles matter because AI can blur the line between preparation and authority. A summary can shape what leaders notice. A classification can change a portfolio view. A recommendation can steer funding. A generated scenario can make one option feel inevitable. The more a system influences decisions, the more explicit the ownership model must be.

The test is simple: if the AI-supported output is wrong, stale, incomplete, biased, or over-relied on, who notices, who corrects it, who explains the impact, and who changes the workflow so the same failure does not repeat? If the answer is vague, the intervention is not ready for meaningful reliance.

What good adoption reporting looks like

Constraint-first adoption reporting should combine activity, workflow, and interpretation. Activity shows whether the tool or system is being used. Workflow measures show whether the bottleneck moved. Interpretation explains whether the change is credible enough to justify the next decision.

A useful monthly view might include the target workflow, current baseline, adoption depth, concentration of use, operational movement, quality sample, exception pattern, cost signal, and decision needed. The report should make clear whether the evidence supports scaling, refinement, a different intervention, or a stop.

The most important row is often the one that says no. No clear baseline. No process owner. No adoption outside enthusiasts. No movement in the workflow metric. No risk control for higher reliance. These are not embarrassing admissions. They are the operating signals that protect the portfolio.

The reporting discipline also helps separate learning from value. A pilot can produce useful learning even when it does not justify scaling. That learning should be recorded: the constraint was misframed, the data was weaker than expected, users did not need the output, the model was too expensive for the use case, or a simpler workflow change would work better. Good governance preserves that learning so the next AI request starts smarter.

Evidence gaps this paper cannot close

This paper does not claim a universal failure rate for all enterprise AI work. The cited sources use different populations, definitions, and research methods. Survey findings, advisory research, and interview-based reports should be read as converging signals, not interchangeable measurements.

It also does not claim that every AI opportunity must produce immediate financial impact. Early exploration can be legitimate when option value is high. The governance question is whether the organization knows it is buying option value, technical learning, workflow evidence, risk reduction, or measurable operating improvement.

The gap this paper can close is managerial. It gives portfolio leaders a first gate that makes AI demand comparable: the constraint, the minimum effective intervention, the owner, and the proof. That standard is practical enough to use before a vendor demo becomes a funded story.

The shift

The next phase of AI adoption will not be won by the organizations with the most pilots. It will be won by the organizations that can decide which work deserves to change, which intervention belongs on the job, and what evidence earns the next level of trust.

That is portfolio work. AI makes execution cheaper and more available. It does not decide which constraint matters. It does not own the risk of relying on an answer. It does not know whether a workflow should be redesigned, simplified, stopped, or left alone.

Start with the constraint. Then choose the tool.

Notes and sources

1. Alex Singla, Alexander Sukharevsky, Bryce Hall, Lareina Yee, Michael Chui, and Tara Balakrishnan, "The state of AI in 2025: Agents, innovation, and transformation," McKinsey & Company, November 5, 2025. Verified July 7, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

2. James Ryseff, "The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI," RAND Corporation, 2024. Verified July 7, 2026. https://www.rand.org/pubs/research_reports/RRA2680-1.html

3. Arun Chandrasekaran, "Why 50% of GenAI Projects Fail - And How to Beat the Odds," Gartner, January 26, 2026. Verified July 7, 2026. https://www.gartner.com/en/articles/genai-project-failure

4. Alexis Krivkovich and Lucia Rahilly, "AI is everywhere. The agentic organization is not yet," The McKinsey Podcast, April 2, 2026. Verified July 7, 2026. https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/ai-is-everywhere-the-agentic-organization-isnt-yet

5. NIST, "AI Risk Management Framework Core," excerpt from AI RMF 1.0, 2023. Verified July 7, 2026. https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

6. Johannes-Tobias Lorenz, Joshan Cherian Abraham, Robert Levin, and Douglas Ziman, "From promise to impact: How companies can measure and realize the full value of AI," McKinsey & Company, April 24, 2026. Verified July 7, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/from-promise-to-impact-how-companies-can-measure-and-realize-the-full-value-of-ai

7. Antoine Montard, Dago Diedrich, and Tanguy Catlin, "Where AI will create value - and where it will not," McKinsey Quarterly, April 29, 2026. Verified July 7, 2026. https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/where-ai-will-create-value-and-where-it-wont

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