13

WHITE PAPER · AI OPERATING GOVERNANCE

AI Deployment Is a Jurisdiction Decision Before It Is a Model Decision

Set a provisional boundary for scope, data, and authority before shortlisting—then confirm it against the candidate

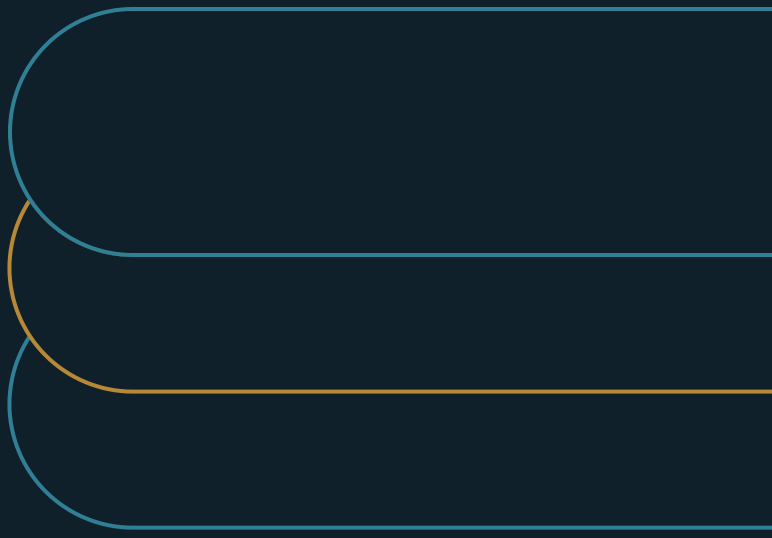

Marco Policani

Enterprise Portfolio · PMO · AI Operating Governance

policani.net · linkedin.com/in/marcpolicani

August 2026

EXECUTIVE SUMMARY

Set the operating boundary first; confirm it against the candidate before commitment.

AI sourcing often starts with model quality, price, and features while the intended use, data path, and action authority remain unsettled. That order lets a candidate's defaults become the enterprise operating model. This paper defines a two-pass operating-jurisdiction review that sets a provisional boundary before shortlisting and confirms it against the selected candidate before commitment.

01 Fix scope, data, authority, responsibility, and approval conditions before a candidate earns a score.

02 Use a two-pass gate: set non-negotiable limits before the shortlist, then test the chosen architecture, contract, and workflow.

03 Keep one decision record with owners, evidence, open conditions, approvers, and triggers that reopen the boundary.

CONTENTS

The decision starts upstream

My operating bias is simple: do not let a product demonstration decide where the enterprise is willing to operate. A model can be impressive and still be wrong for the boundary the business needs. Set the provisional boundary first. Name the work in scope and the data allowed to enter and move. Set the system's maximum authority. Name the accountable owner and the person who can approve an exception.

That boundary is a jurisdiction in the practical sense. It includes law, but it is broader than a legal memo. It is the combined field of external obligations and internal commitments that will govern the deployment after the sales meeting ends. Regulation may define duties. Contracts allocate responsibilities. Security architecture constrains access. Data policy limits movement and reuse. The operating model decides who may rely on an output and who may authorize an action.

Model-first sourcing reverses that order. A team compares quality, speed, and price. It reviews context windows, integrations, and vendor road maps. A favorite emerges. Only then does the organization ask where prompts are processed and whether inputs are retained. It asks which people are affected, what notices are required, and whether the workflow crosses a boundary the business meant to keep. By then, the preferred option has an internal constituency. A boundary question feels like an obstacle to a decision already made.

The better sequence is a two-pass operating decision, not legal review and technology review sitting in separate queues. Before shortlisting, set a provisional boundary that any candidate must fit. Before commitment, test that boundary against the proposed model, architecture, contract, and workflow. A candidate that changes the boundary returns the choice to governance rather than leaving it inside procurement.

Jurisdiction is more than geography

The word jurisdiction usually points to countries, regulators, and courts. Those matter, but a deployment can cross important boundaries without crossing a national border. A customer record can move from an approved business system into a third-party inference service. An assistant can shift from drafting to recommending. A recommendation can become an automated action. A local tool can acquire a cloud dependency through an update. A business unit can extend a bounded workflow to a population the original review never considered.

An operating-jurisdiction review therefore asks three classes of question at once. Scope questions establish who, where, and which work is affected. Data questions establish what information enters, where it travels, how it is used, and when it is removed. Authority questions define what the system may do: retrieve, draft, recommend, decide, or act. They also name who can approve, override, suspend, and answer for each step.

The legal layer sits inside that review. Regulation (EU) 2024/1689 assigns duties through defined roles such as provider and deployer, through intended purpose and use, and through categories of system and context. Headquarters alone does not answer those questions. Neither does the vendor’s label for its product. The relevant facts come from the planned workflow, the people affected, the organization’s role, and the way the system is placed into service. Qualified counsel determines the legal conclusion; the operating team must supply a reliable fact pattern. [2]

The business layer can be stricter than the legal floor. An organization might keep certain records inside a controlled environment while barring automated ranking from a class of personnel decision. Another policy might let a system draft an external message but never send it. Those are operating choices about risk tolerance, reputation, customer promise, and accountability. They belong in the deployment decision even when no statute dictates the exact answer.

This paper is not legal advice and does not offer a territorial-scope or compliance test. Its claim is narrower. Make the facts and decision rights behind a deployment boundary explicit before a model or vendor defines them by default.

The August 2026 milestone exposes the sequencing problem

The European Union's AI rulebook makes the timing visible because its obligations do not arrive as one undifferentiated event. The original AI Act set 2 August 2026 as its general application date. The European Commission now describes 2 August 2026 as the point from which the AI Office and Member State authorities implement, supervise, and enforce the Act. It also notes that rules for general-purpose AI models became effective in August 2025. [1] The distinction matters: "the AI Act applies" is not a sufficient operating instruction.

The Digital Omnibus on AI is no longer a proposal awaiting publication. Regulation (EU) 2026/1744 was published in the Official Journal on 24 July 2026 and entered into force on 27 July. It moved most of Chapter III, Sections 1, 2, and 3 to later dates. For systems classified under Article 6(2) and Annex III, the new date is 2 December 2027. For systems classified under Article 6(1) and Annex I, it is 2 August 2028. Article 6(5) is excluded from that move. The regulation also gives some generative systems placed on the market before 2 August 2026 until 2 December 2026 to meet Article 50(2) marking duties. [3]

Those dates correct an easy but consequential shorthand. The 2 August 2026 milestone is real, but it is not the day every obligation began. GPAI duties started earlier. Some transparency duties have a transition. Major high-risk requirements now start later. A deployment review that carries only a single "AI Act date" can be accurate at headline level and wrong at control level.

The point is to stop the portfolio from treating regulatory status as a yes-or-no field attached to a vendor, without asking every program team to interpret legislation. The intake record should identify the planned use, actors, affected population, system role, data path, and deployment locations. Legal and compliance owners can then map those facts to current obligations. When dates or classifications change, the organization can find the deployments whose boundary is affected without repeating discovery from scratch.

The same record serves jurisdictions beyond the EU. It does not assume that one regime governs the world. It creates a stable set of operating facts that different legal, regulatory, contractual, and policy

owners can assess. That is more durable than embedding one regulation's current checklist inside a purchasing form.

Why model-first selection creates boundary debt

A model choice is rarely just a model choice. It tends to bring a hosting pattern, identity model, data path, monitoring stack, update cadence, supplier chain, contract, and product defaults. Each choice may be reversible in theory. Together, they create switching costs. Integrations are built. Users are trained. Sponsors announce a direction. Data mappings and security exceptions are approved. The organization begins paying for its first assumption before it has named the assumption.

That is boundary debt: the work required to recover a scope, data, or authority limit that should have been explicit at intake. It appears in several ways. A vendor's standard retention period becomes enterprise policy by convenience. A "copilot" starts transactions because the integration supports it. A global user group is enabled before regional differences are reviewed. A pilot uses production-like data because the environment made it easy.

Boundary debt can arise without vendor misconduct. Suppliers design products for repeatable delivery. Their defaults express their architecture, market, and risk allocation. The buyer decides whether to accept, configure, contract around, or reject those defaults. CISA's cloud architecture guidance describes a shared model for security. The cloud provider and customer each hold part of the responsibility. Neither side receives it all. [6] The exact split changes by service and agreement. A usable responsibility map assigns the edges that vendor and enterprise slogans leave open.

NIST's AI Risk Management Framework is useful here because it begins with organizational context rather than model rankings. It is designed for organizations that develop, deploy, or use AI across sectors and use cases. [4] Its Core calls for a clear intended purpose, business context, and risk tolerance. It also calls for relevant laws, roles, and third-party risks to be understood and recorded. [7] The Generative AI Profile extends that frame across design, development, use, and evaluation. [5] Those conditions make a procurement score meaningful.

The operating-jurisdiction review

The review can be run as a compact intake gate. It does not need to become a universal committee or a months-long pre-clearance process. Its weight should follow consequence and reversibility. A private drafting experiment with synthetic data can use a light review. A system that ranks applicants, handles regulated records, or starts payments requires more. Both need a named boundary.

EXHIBIT 1

Five decisions set the deployment boundary

01	02	03	04
Scope	**Data**	**Authority**	**Responsibility**
Fix the use, exclusions, users, affected people, locations, and intended outcome.	Approve data classes, movement, retention, training use, access, and deletion.	Set the strongest allowed action, required human decision, and stop point.	Assign enterprise and supplier duties for change, evidence, incidents, and halt.

05

Approval

Record outcome, conditions, approvers, review date, and re-entry triggers.

The artifact is built around five decisions. They are stated as decisions rather than information requests. A form can collect facts without resolving who accepts the consequence. Each decision needs an owner and a record of the evidence reviewed. The record marks each item as fact, judgment, or open question. It names an owner and date for any open condition. It also shows whether the decision is in pass one or pass two.

1. Fix the operating scope

Name the business process, the users, the people affected, the deployment locations, and the intended outcome. Define what is outside scope with equal care. “Enterprise assistant” is not a scope. “Drafts internal service responses for a named support group, using an approved knowledge set, with a human sending every response” is a scope that can be tested.

Scope also includes the organization’s role. Is it buying a finished service or changing a system? Is it embedding a model in its own offering or placing a capability in front of customers? The legal label belongs to counsel, but the operating facts belong to the team. Record them before a vendor presentation encourages the organization to describe the use in the supplier’s language.

2. Fix the data boundary

List the data classes the workflow may receive, retrieve, generate, and expose. For each class, record its origin and permitted use. Record storage, retention, deletion, training use, transfer, logging, and

access. Include derived data and outputs. A rule that protects the prompt but ignores embeddings, logs, feedback, cached responses, or tool results is not a complete boundary.

The review should state prohibited inputs in terms users can recognize. “No sensitive data” is too abstract. The operating record should point to the organization’s actual classes and examples. It should name the approved source systems and the response to an input outside the boundary. Prevention, warning, quarantine, deletion, and incident response are different controls. The team must choose among them.

3. Fix the authority boundary

Describe the strongest action the system may take. Retrieval, summarization, drafting, recommendation, ranking, decision, and execution are not interchangeable. For each action, name the human role that reviews or authorizes it, the conditions under which review is required, and the point at which the system must stop. If no one can suspend the workflow quickly, the deployment has an approval path but no recovery path.

Human review is not a generic checkbox. A reviewer needs time and access to the relevant evidence. The reviewer also needs authority to reject the output and a clear standard. An approval step that people cannot perform is workflow decoration. The record should specify what the reviewer sees and what decision the reviewer makes. It should state which evidence is kept and how doubt or disagreement is escalated.

4. Fix the responsibility split

Map responsibilities across the business, technology, data, security, privacy, compliance, procurement, legal, and supplier roles. Then assign each action. Who monitors service changes, receives incident notices, keeps logs, reviews major updates, answers inquiries, and halts use? Shared responsibility becomes governable when each shared edge has an owner.

Supplier evidence belongs in this step. It may cover architecture, data terms, subprocessors, security controls, evaluation, change notices, incidents, continuity, and exit support. The purpose is to obtain the evidence the boundary requires, not every document from every vendor. A low-consequence use may need a small evidence set. A consequential deployment earns deeper proof and stronger contract rights.

5. Fix the approval and re-entry conditions

State who can approve the provisional boundary and who can approve the final candidate. Then name the changes that reopen the decision. A new user group, geography, data class, or automated action may trigger review. So may a new model family, hosting region, supplier, integration, or major capability. Without triggers, the deployment can drift into a different jurisdiction while retaining the original approval.

Approval should produce one of four outcomes: proceed within the boundary; proceed with named conditions; return for evidence or redesign; or stop. “Approved” by itself is weak because it hides the commitments attached to the decision. Conditions need an owner and a latest useful date. A condition that remains open after production access is usually an unrecorded transfer of risk.

The decision record

The review produces a decision record that is short enough to use and specific enough to inspect. Its first block fixes the use case, exclusions, users, affected people, and locations. It also notes legal questions routed for advice. The second block fixes allowed and prohibited data, processing, retention, transfer, and maximum system authority. It names the human decision that cannot be delegated. The last block assigns responsibilities, evidence, open conditions, approvers, re-entry triggers, and the next review date.

A useful record also preserves the alternatives that were rejected. The team may have considered a local deployment, a managed service, a smaller model, a non-AI workflow change, or waiting for a control. Record why the selected path fit the boundary better. This prevents the decision from being rewritten later as an inevitable technology choice. NIST includes viable non-AI alternatives in risk treatment. The portfolio should do the same when the boundary makes an AI option expensive or brittle.

The record separates facts, judgments, and open questions. “Inputs are not retained under the proposed contract” is a fact to verify. “The residual exposure is acceptable for this use” is a judgment with an owner. “Whether the organization is a provider for this configuration” is a legal question for qualified review. Mixing those statements in one status field makes each less reliable.

This separation also makes later change visible. A new contract term changes a fact. A larger user group may change the judgment. New law may change the question and its answer. The team can reopen the affected part of the record. It need not treat the whole deployment as always approved or always provisional.

A two-pass gate, not a one-time declaration

The strongest objection to jurisdiction-first governance is correct: the model and architecture can change the boundary. A system’s capabilities, integration method, update behavior, supplier chain, and contractual terms matter. It would be false precision to declare a final jurisdiction before those facts exist.

That is why the review has two passes. Pass one occurs before the shortlist. It fixes the intended use, excluded uses, data classes, maximum authority, accountable owner, and firm operating limits.

Candidates that cannot fit are not scored. Pass two occurs before commitment. It tests the chosen candidate against the provisional record. It adds legal and technical findings, resolves conditions, and names re-entry triggers.

The second pass may change the first. That revision is due diligence working as intended. The control is that the change is explicit. A candidate may require a broader data path or stronger authority. If the sponsor still wants it, that change returns as a boundary decision with its consequences visible.

This sequence also protects speed. Teams stop spending time on options the organization cannot deploy. Legal, security, data, and procurement reviewers receive a stable fact pattern instead of a generic request to “review the AI.” Vendors receive sharper questions. The portfolio sees where several use cases depend on the same unresolved boundary and can fund a shared control instead of negotiating one exception at a time.

Who owns the review

The accountable business owner owns the intended use and the decision to seek value within a stated boundary. Technology and architecture owners establish how the workflow operates. Data owners approve permitted data use. Security, privacy, compliance, procurement, and legal counsel supply decisions within their mandates. None of those functions should become the default owner of the business decision simply because the system contains AI.

EXHIBIT 2

Ownership follows the decisions the review must close

Business owner

Own the intended use, value case, consequence, and request to operate inside a boundary.

Control owners

Close the data, security, privacy, compliance, legal, and contract decisions in scope.

Technology & supplier

Provide architecture, evidence, change notice, incident, continuity, and exit facts.

Portfolio governance

Own the gate, decision record, escalation path, and cross-portfolio constraint view.

Portfolio or AI-governance leadership owns the gate design and the record standard. Its job is to show which deployments cross similar boundaries and where evidence is missing. The record should also

reveal growing exceptions and shared capabilities that would cut repeated work. The gate should not turn every use case into the same review. Proportionality is part of the design.

One accountable owner can complete a small review with pre-approved patterns. A material deployment may require a cross-functional decision. Escalation does not depend on whether the tool uses a fashionable model. It depends on consequence, reversibility, data exposure, affected population, and delegated authority.

What the evidence does—and does not—establish

The legal sources establish current EU roles, dates, and categories. NIST offers a voluntary risk frame for many sectors. It covers context, documented responsibility, lifecycle risk, and third parties. CISA sets out shared responsibility for cloud security. Together, these sources support defining context and responsibility before deployment. They do not validate this exact five-decision review. Nor do they prove that it reduces losses or set the right boundary for a given organization.

EXHIBIT 3

The sources bound the review; they do not prove its return

- **EU law**
  Establishes current roles, application dates, transparency duties, and high-risk timing.
- **NIST guidance**
  Supports context, ownership, risk tolerance, lifecycle controls, and third-party review.
- **CISA guidance**
  Supports an explicit split of cloud security responsibility between provider and customer.
- **Operating synthesis**
  Defines the five decisions, two-pass gate, record, and re-entry triggers used here.
- **Evidence gap**
  No causal result, universal threshold, or legal conclusion is claimed for this review.

The evidence gap matters. No cited base rate shows how often model-first selection causes boundary debt. This paper also has no comparative study showing that a jurisdiction-first gate improves delivery outcomes. The review is an operating control drawn from portfolio, risk, and decision-rights practice.

The practical test is whether it exposes a material boundary before commitment and routes the right question to the right owner. A later reviewer should also be able to reconstruct the decision.

Exploration still has a place

Jurisdiction-first does not require the organization to know every boundary before it learns what a model can do. Some questions require a test. A team might compare a smaller model on the task or test whether retrieval improves answers. It may also observe whether reviewers can keep pace and whether a proposed control is usable. Blocking all exploration until the final review would hide those facts and turn governance into speculation.

The answer is a narrow learning jurisdiction. Use synthetic or approved non-production data. Limit participants and integrations. Disable external actions. Set an end date and a deletion plan. State which conclusions the experiment can and cannot support. The experiment may inform the final boundary. It may not quietly become production because it gained users or produced a persuasive demonstration.

This distinction protects both learning and control. A sandbox is a small jurisdiction designed for uncertainty. Its lighter controls are justified by bounded consequence and reversibility. Real customer records, broader access, a production integration, or an action path move the work into a different jurisdiction and trigger re-entry.

The pass-one entry screen

A sponsor should be able to answer seven questions before the sourcing team compares models. This fast screen opens the five-decision record. It tests whether the use case is clear enough to enter the market. Reviewers then resolve and approve the detailed scope, data, authority, responsibility, and re-entry decisions.

- What exact work will change, and what work will remain outside the deployment?
- Which people will use the system, and which people can be materially affected by its output or action?
- Which data classes are required, prohibited, or still unresolved?
- What is the strongest action the system may take, and which human decision cannot be delegated?
- Which deployment locations, business entities, customer commitments, and external obligations require review?
- Which supplier facts or contractual rights must be true for this use to fit the boundary?
- Who can approve, suspend, and reopen the decision when the use, data, model, or service changes?

A blank answer does not always mean stop. It means the blank has an owner and a due point before the relevant commitment. A team can test task quality before it fixes a production hosting region. That test must use approved data and take no external action. The team cannot sign a contract that fixes the hosting pattern while calling the region question open.

This is also where portfolio discipline improves the economics of governance. Several use cases may stall for the same reason: a missing data path, logging service, contract clause, or review capacity. The portfolio has then found a shared constraint and can fund it on purpose. Without the records, each project appears to have a unique delay. The enterprise pays repeatedly for local workarounds.

The boundary must survive change

Deployment is not the end of the jurisdiction decision. Models are updated, services add features, contracts renew, integrations deepen, and users find new applications. A review that captures only the launch configuration expires as soon as the system changes. The decision record therefore needs a monitoring owner and a small set of observable triggers.

The owner does not need to treat every vendor release note as a new project. The owner watches for changes that can alter the accepted boundary. Product triggers include a new model family, tool use, or a changed review design. Data and supplier triggers include new retention terms, training terms, a subprocessor, a region, or a source. Scope triggers include a broader population or a new automated action. An incident that disproves the assumed responsibility split is also a trigger.

Those triggers turn monitoring into decision support. They identify when the old approval no longer describes the operating system. The response may be confirmation, a control adjustment, a temporary restriction, or a stop. What matters is that continued use becomes a conscious decision when the boundary moves.

Choose inside the boundary

Once the operating jurisdiction is explicit, model selection becomes more useful. Teams can compare quality on the real task, total operating cost, latency, and integration effort. They can also compare evidence, controls, and supplier terms inside a chosen boundary. A candidate that misses a firm data or authority limit falls outside the decision.

The order is the point. Define the work before the tool. Define the data path before the contract. Define the system’s authority before the integration. Define who can approve, stop, and reopen the decision before production. Then choose the model that fits.

AI deployment is not governed by the model with the highest score. It is governed by the boundary the organization is prepared to defend. That boundary must hold when the workflow changes, a regulator

asks, a customer objects, or a supplier updates the service. It must also hold when an incident exposes who owned the decision.

Notes and sources

[1] European Commission, AI Act, Shaping Europe's digital future, updated 3 August 2026.

[2] European Parliament and Council of the European Union, Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence, Official Journal of the European Union, 12 July 2024.

[3] European Parliament and Council of the European Union, Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), Official Journal of the European Union, 24 July 2026.

[4] Elham Tabassi, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, National Institute of Standards and Technology, 26 January 2023.

[5] Chloe Autio, Reva Schwartz, Jesse Dunietz, Shomik Jain, Martin Stanley, Elham Tabassi, Patrick Hall, and Kamie Roberts, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, National Institute of Standards and Technology, 26 July 2024.

[6] Cybersecurity and Infrastructure Security Agency, United States Digital Service, and Federal Risk and Authorization Management Program, Cloud Security Technical Reference Architecture, Version 2.0, June 2022.

[7] National Institute of Standards and Technology, AI RMF Core, Artificial Intelligence Resource Center, excerpt from Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023.

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