OMNITEK CONSULTING

Governance for AI Pilots

WHITE PAPER

Sponsoring, Tracking, and Sunsetting AI Initiatives Before They Become Shadow IT

Abstract

Eighty percent of healthcare AI projects fail to scale beyond the pilot phase, and the reasons are rarely about the technology itself. Fragmented data environments, governance that arrives after a pilot is already running rather than before it starts, and the absence of any defined process for moving from experiment to production leave promising pilots stranded indefinitely, neither shut down nor scaled, quietly consuming budget and, worse, quietly becoming unsanctioned production tools nobody formally approved. Regulators are not waiting for the industry to sort this out informally. The FDA's January 2025 draft guidance on AI in drug development lays out a seven-step, risk-based credibility framework that any AI model touching a regulatory decision now has to satisfy. This white paper applies established program governance best practices to sponsoring, tracking, and, when warranted, deliberately sunsetting AI pilots before they drift into ungoverned use.

The Pilot That Never Officially Started or Stopped

Most AI initiatives inside life sciences organizations begin the same way: a team with a real problem finds a promising tool, runs an informal proof of concept, and gets encouraging early results. What happens next is where most of these initiatives quietly stall. Eighty percent of healthcare AI projects fail to scale beyond the pilot phase, not because the underlying technology failed, but because the organization around it was never built to move a pilot from experiment to production.[i] Fragmented data environments across research, clinical, and commercial functions limit the ability to see how a pilot's data relationships actually hold up at scale, and when governance oversight only shows up after a pilot is already running, it functions as a roadblock rather than a structure the pilot was designed around from the start.[i]

The failure mode that should worry PMO and governance leaders most is not the pilot that gets formally shut down. It is the pilot that never gets a formal decision either way, that keeps running in a business unit because it produced early value and nobody assigned anyone to decide whether it should continue, scale, or stop. That kind of pilot does not disappear. It becomes exactly the shadow IT problem governance was supposed to prevent, a tool making real decisions inside the business with no formal risk assessment, no defined ownership, and no plan for what happens when the person who built it moves to a different role.

Why Regulators Are Not Waiting for Informal Governance to Catch Up

The FDA's January 2025 draft guidance on artificial intelligence in drug and biological product development made clear that regulators are done treating AI governance as optional or informal, at least for any model whose output feeds a regulatory decision about safety, efficacy, or quality.[ii] The guidance centers on establishing model credibility, meaning documented trust in an AI model's performance within a specific, defined context of use, through a seven-step process: define the regulatory question, specify the model's context of use, assess the risk that context carries, build a credibility plan proportionate to that risk, execute the plan, document results and any deviations, and determine whether the model actually serves its intended purpose.[ii]

That framework was written for models feeding regulatory submissions specifically, but its underlying logic, define the use case, assess its risk, build oversight proportionate to that risk, document everything, applies just as well to the broader universe of AI pilots running inside a life sciences organization that never touch a regulatory filing directly. A pilot summarizing internal literature for medical affairs carries a very different risk profile than one flagging high-risk patients for a clinical trial, and treating both under the same informal, ungoverned pilot process, or the same heavy formal process, misses the proportionality the FDA itself is now building into its own regulatory expectations.

Applying Best Practices: A Governance Model for AI Pilots

A foundational best practice in PMO governance treats Governance and Decision Rights as a core maturity dimension, and AI pilot governance is one of the fastest-growing places that dimension gets tested. Applying that discipline to a specific AI pilot model gives PMOs a practical way to sponsor promising initiatives without losing track of them. Five components make up that model.

1. **A Defined Intake Gate.** No pilot begins without a documented owner, a stated business problem, and an explicit context of use, borrowing directly from the FDA's own framing, before any tool gets connected to real data or a real workflow.
2. **Risk-Proportionate Oversight.** The rigor applied to a pilot's governance scales with the risk of its context of use, a low-stakes internal efficiency pilot warrants lighter oversight than one touching patient safety, clinical decisions, or regulated data, mirroring the proportionality at the center of the FDA's credibility framework.[ii]
3. **A Time-Boxed Pilot Clock.** Every pilot gets an explicit review date at launch, not an open-ended runway, forcing a deliberate decision, scale, extend with justification, or sunset, rather than allowing indefinite drift into informal production use.
4. **A Real Path to Production.** Pilots that succeed have a defined, resourced path to move from proof of concept into a properly governed production tool, closing the gap where even genuinely successful pilots stall simply because no process existed to formalize them.[i]
5. **A Documented Sunset Process.** Pilots that do not warrant scaling get a formal, documented decommissioning, including removing data access and closing out any informal workflow dependency that built up around the tool, rather than being quietly abandoned while still technically live.

The goal is not to slow down experimentation. It is to make sure every pilot has a name attached to it, a clock running against it, and a defined outcome waiting at the end of that clock, so promising ideas either scale properly or get closed out cleanly, and nothing lingers in the ungoverned middle.

What AI Pilot Governance Buys, and Where It Gets Hard

Advantages

- **Fewer shadow AI tools:** A defined intake gate and pilot clock close the gap where ungoverned tools quietly become part of real business workflows without ever passing through a formal risk assessment.
- **Faster path from pilot to value:** A real, resourced production pathway addresses the specific failure mode behind the eighty percent scaling failure rate: promising pilots stalling for lack of any process to formalize them.[i]
- **Regulatory readiness:** Organizations already applying context-of-use and risk-proportionate thinking to internal AI pilots are better positioned when a pilot's output eventually needs to satisfy the FDA's credibility framework for a regulatory submission.[ii]
- **Clean, documented accountability:** When a pilot's outcome is ever questioned, whether by an auditor, an inspector, or an internal risk review, a documented intake, oversight, and sunset record provides a defensible answer instead of an informal, reconstructed one.

Risks and Barriers

- **Governance overreach:** Applying heavy, uniform oversight to every pilot regardless of risk recreates the exact problem the framework is meant to solve, discouraging the low-risk experimentation that should move quickly.
- **Sunset resistance:** Teams that have built real workflow dependency on a pilot resist a formal sunset decision even when the pilot genuinely should not scale, and that resistance needs a clear, principled process to work through rather than an indefinite extension.

- **Intake gate bottlenecks:** An intake process with too few resources to review pilots promptly becomes its own obstacle, pushing teams back toward informal, ungoverned experimentation to avoid the delay.
- **Cross-functional ownership gaps:** AI pilot governance sits at the intersection of IT, quality, legal, and the business function running the pilot, and without clear joint ownership, accountability for any single pilot can fall through the cracks between functions.

Case in Point: The Literature Surveillance Tool Nobody Owned

A medical affairs team at a mid-size biopharma piloted an AI tool to surface relevant new literature for a specific therapeutic area, built informally by a data-savvy team member using a departmental budget with no formal IT or governance review. The pilot worked well enough that within a year, three additional therapeutic area teams had adopted informal versions of the same workflow, each with its own configuration, none registered anywhere as an active AI initiative, and none owned by anyone accountable for its data sourcing or output accuracy.

The gap surfaced when the original tool's creator left the company and none of the four teams using variants of the workflow had documentation sufficient to maintain or validate what they were actually relying on. A subsequent internal review found no evidence of any output errors having caused harm, but it also found no way to rule that out with confidence, precisely because nothing about the tool's context of use, risk profile, or ongoing performance had ever been formally assessed. Retrofitting governance onto four years of informal use took considerably longer than building it in from the first pilot would have, and left the organization unable to answer, with real confidence, how reliable the tool's output had actually been the whole time it was in use.

Recommendations for PMO and Governance Leaders

1. **Build the intake gate before the next pilot request arrives.** Define the minimum information required to launch any AI pilot, owner, business problem, context of use, before the next team asks to start one informally.
2. **Scale oversight to risk, not to enthusiasm.** Resist applying uniform governance rigor regardless of a pilot's actual risk profile. A low-stakes internal tool and a clinical-decision-support pilot warrant genuinely different levels of scrutiny.
3. **Give every pilot a review date at launch.** No pilot should have an open-ended runway. Build the scale-extend-sunset decision point into the pilot's charter from day one.
4. **Resource the path from pilot to production.** A successful pilot with nowhere to go is a wasted investment. Fund and staff the formalization process as deliberately as the pilot itself.
5. **Make sunsetting routine, not punitive.** Frame a pilot that does not scale as a normal, expected outcome, not a failure, so teams report honestly rather than quietly keeping an unsanctioned tool alive.
6. **Assign joint ownership across IT, quality, legal, and the business.** Name specific accountable owners from each function for the pilot governance process itself, not just for individual pilots, so cross-functional gaps do not become nobody's responsibility.

Measuring Success: KPIs for AI Pilot Governance

- **Intake gate coverage:** the percentage of active AI pilots that passed through the formal intake process versus those discovered informally after the fact, tracked toward full coverage.
- **Pilot-to-production conversion rate:** the share of pilots that reach a formal scale decision within their reviewed timeline, addressing the industry's broader eighty percent stall rate directly.[i]
- **Time to sunset decision:** how long a pilot runs past its scheduled review date before a formal scale, extend, or sunset decision is made, tracked toward zero drift.
- **Shadow AI discovery rate:** the number of AI tools found in active use that never passed through governance, ideally trending toward zero as the intake gate becomes the default path.
- **Risk-proportionate oversight accuracy:** whether pilots receiving the most rigorous oversight actually correspond to the highest-risk contexts of use, confirming the proportionality principle is being applied in practice, not just on paper.

Conclusion

The eighty percent of AI pilots that never scale are not, for the most part, technology failures. They are governance failures, pilots launched without a clear owner, a defined context of use, or any process for reaching a real decision about their future, left instead to either quietly die or quietly become the exact kind of ungoverned shadow tool a well-run PMO exists to prevent. Regulators have already signaled where this is heading, building risk-proportionate, documented credibility assessment directly into how AI touching regulatory decisions gets evaluated.

Applying disciplined best practices to a deliberate AI pilot governance practice does not mean slowing innovation down with heavy process. It means giving every pilot the same basic structure regulators are already requiring for the highest-stakes AI use cases: a defined purpose, oversight proportionate to actual risk, and a real decision point instead of an indefinite runway. Organizations that build that structure now will not just avoid the shadow AI problem. They will already be positioned for the governance regulators are going to expect from them regardless.

About OmniTek & Contact Information

OmniTek Consulting is a trusted partner for organizations seeking to enhance governance, operational agility, and decision-making. We specialize in modernizing business processes with a blend of human-centered design, technology expertise, and hands-on change management.

For PMO and governance leaders trying to get ahead of AI pilots before they become ungoverned shadow tools, OmniTek applies proven best practices to build a practical governance model: a defined intake gate, risk-proportionate oversight, and a real path to either production or a clean sunset. Whether your organization is standing up AI governance for the first time or trying to formalize pilots that have already been running informally for years, OmniTek can help you build a process that lets innovation move quickly without losing track of what is actually running inside your business.

Contact us today to explore how OmniTek can support your transformation journey and drive lasting results.

McLean, VA | Serving clients nationwide
info@omnitekconsulting.com
omnitekconsulting.com

Click here to book a 15-minute call to see how OmniTek can help your team put these insights into practice.

Endnotes

[i] HIT Consultant, "The 'Pilot Purgatory': Why 80% of Pharma AI Projects Fail (And How to Fix It)," February 2026, available at https://hitconsultant.net/2026/02/16/cloudera-life-sciences-ai-scalability-pilot-failure-roi/

[ii] Goodwin, "FDA Publishes Draft Guidance on Use of Artificial Intelligence in the Development of Drugs and Biological Products," January 2025, available at https://www.goodwinlaw.com/en/insights/publications/2025/01/alerts-lifesciences-aiml-fda-publishes-its-first-draft-guidance