incedo

Whitepaper

# The Foundation Tax

Why launch readiness breaks at portfolio cadence, and what to build first

# Executive Summary

Launch misses have persisted for two decades while the causes have been published, workshopped and largely agreed. When known problems survive known solutions for that long, the constraint is not knowledge. It is capacity.

We call the capacity problem the foundation tax: the delay, rework and avoidable spending that recur every time meaning, assets, rights metadata and workflow state are not shared across launches. At one launch a year it is invisible. Across a portfolio it compounds, and it lands hardest on the two supply chains that decide whether a launch is ready, which are data readiness and content readiness.

Two arguments follow, and both cut against current practice. Agents should assemble the evidence for consequential decisions and stop there, because the context that governs those decisions is usually not in the data. And the foundation has to be sequenced before the applications, built through two instrumented workflows rather than as a platform program, or every launch pays for it again.

## The renewal agent that knew too little

Not long ago our team proposed what looked like an obvious win inside a large biopharma commercial data organization. Data subscriptions come up for renewal on a calendar. Utilization is logged. Why not build an agent that reads the usage signals and recommends which contracts to renew and which to drop?

The operators who ran that portfolio pushed back, and they were right. Most of their spending sat in core sales, claims and prescription assets that stayed necessary for as long as a product was in market. A low usage reading on one of those assets said nothing about whether the company could do without it. Dependencies, therapeutic context and stakeholder commitments never appeared in the logs. An agent acting on utilization alone would have been confidently wrong about the decisions that mattered most.

So the design changed. The system now assembles the record a human decision needs: usage, cost, contract terms, therapeutic area, downstream dependencies and the stakeholders involved. Client operators preferred the narrower concept because it removed coordination work and left the consequential decision with the accountable owner. That preference shaped the design. It does not by itself demonstrate adoption, time saved or money saved, and the distinction between the two runs through everything below.

**We think the industry has this backwards. The prevailing bet is that agents will get good enough to make these calls. The harder and more useful problem is that the context governing them is not written down anywhere an agent can read. Automate evidence assembly first. Judgment can wait, and the consequential call itself should wait until an accountable owner signs.**

## What the foundation tax is, and how you would know we are wrong

The same avoidable work appears in every large biopharma environment our teams have worked in. A team evaluates or acquires data without a complete view of what the portfolio already owns. A vendor assessment restarts because the previous decision and its rationale were never made reusable. The same business term means different things in the request system, the procurement system and the reporting layer. Content teams repeat review work on claims and modules that already have an approved source, even though each newly assembled communication still needs its own review.

That is the foundation tax. It is not an accounting category and we have not measured its prevalence across the industry. It is an operating hypothesis, and a hypothesis has to be falsifiable to be worth anything.

So here is the test. If a shared foundation reduces the tax, an organization should see shorter request-to-ready times, fewer duplicate evaluations and acquisitions, less content rework, and more readiness gates met by their decision dates. Control outcomes should be measured against actual exposure, not against the absence of observed incidents. If those numbers do not move, the foundation is architecture and nothing more. We would rather publish the kill criteria than defend the idea after the fact.

# Why cadence changes the math

The economics are unforgiving and well documented. Deloitte's sixteenth pharmaceutical innovation analysis, published in May 2026, puts average discovery-to-launch development cost at 2.671 billion dollars in 2025. Health Advances, studying fifteen successful launches by single-product public companies, estimated 475 to 830 million dollars of commercialization spending across a six-year window. On outcomes, McKinsey found roughly two thirds of 210 new molecular entities launched between 2003 and 2009 missed first-year consensus expectations, and Deloitte found 34 percent of 284 United States launches between 2012 and 2021 generated no more than 80 percent of expected launch-year sales. Different cohorts and thresholds, so they do not form a trend line.

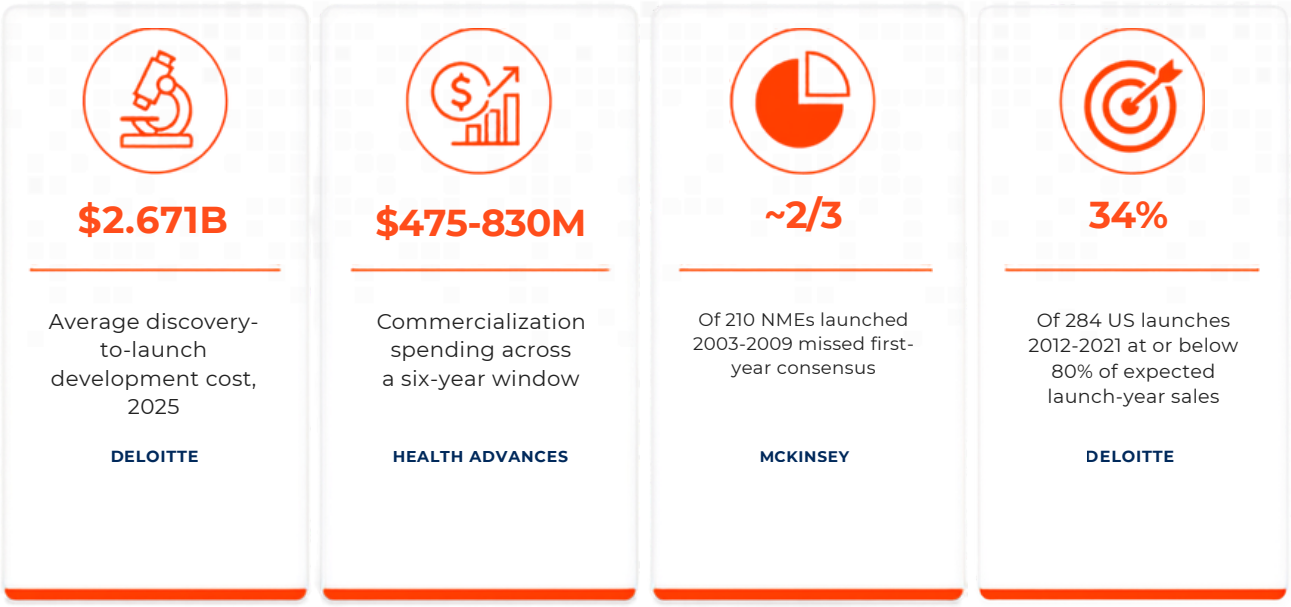


Deloitte's finding on trajectory is the one that should worry people. Of the launches that missed in year one, only 26 percent recovered in either of the next two years. A weak first year is close to irreversible, which means readiness has to be right before the window opens rather than corrected inside it.

What has changed is not approval volume. FDA approved 50, 37, 55, 50 and 46 novel drugs across 2021 through 2025, which describes system-level flow and says nothing about any one company's calendar. What has changed is overlap. Product introductions, indication expansions, formulation changes and market entries now land on the same calendar and draw on the same data acquisition, content review and governance capacity. IQVIA projects United States loss-of-exclusivity exposure above 90 billion dollars at estimated net manufacturer prices between 2025 and 2029, which is the pressure forcing the replacement pipeline through a fixed set of shared resources.

At low cadence, experienced teams paper over fragmented systems with manual coordination and institutional memory. As events stack up, that stops working and delays propagate across the portfolio. The measure worth adopting is portfolio launch throughput: the share of in-scope events that pass every pre-specified readiness gate by the decision date each gate serves, inside the quality, rights, privacy and compliance guardrails. It is a management measure rather than an industry standard, and its value is that it forces an organization to name which events count and which constraints speed is never allowed to violate.

## Data readiness: a design case

Launch data requests move through request portals, procurement systems, data platforms, spreadsheets, contracts and email. In several environments nobody could see, in one place, how a request connected to existing assets, permitted uses, prior vendor assessments, acquisition status and the launch decision the data was meant to serve.

Two risks follow, and both are countable. A team buys coverage the portfolio already holds under usable rights. And usable data arrives after the access, segmentation or measurement decision it was supposed to inform.

Our teams have piloted agentic AI orchestration against this. Task-specific agents handle intake structuring, portfolio overlap detection, contract and rights retrieval, and reconciliation, coordinated by an orchestrator that shares context between them and routes open questions to accountable owners. Underneath sits a semantic layer supplying agreed definitions for terms such as usable, covered population, launch critical and approved purpose. That layer is what makes agent output trustworthy instead of merely fluent, and it is the part most organizations skip. These pilots demonstrate technical feasibility. They do not demonstrate business impact, and an overlap flag is not a confirmed duplicate until an accountable owner validates that the existing coverage would have served the request. The outcomes that count are specific: eligible requests processed, flags confirmed or rejected, evaluations stopped before vendor engagement, request-to-ready time against a baseline, adoption by the people who own the work, and exceptions per hundred requests. Until numbers like those can be disclosed this remains a design case, and we would rather say so than imply otherwise.

## Content readiness: where the constraint moves next

Neither establishes a universal constraint, which is why an organization should measure its own queue instead of assuming saturation: days before review, days in review, review rounds, first-pass acceptance, rework and critical deviations.

This is where AI-driven dynamic content generation is being applied now. In a second large pharma engagement our teams are building agents that retrieve approved sources, assemble modular content for a given channel and audience, check references, attach metadata and route the result for review. The architecture bet is identical to the data side. Generation quality depends on whether approved claims, evidence currency and channel rules are governed underneath, or inferred at the point of drafting. Ungoverned, a generation engine produces more work for the review queue, not less.

None of this removes accountable review of the assembled communication. FDA expects prescription drug promotion to be truthful, balanced and accurately communicated, and covered postmarketing promotional materials must be submitted at initial dissemination. Prior approval of a source module does not settle a new communication's audience, channel, surrounding text, evidence currency or fair balance. The workable objective is compliant velocity: design the review workload down through modularity and reuse, and keep the approval decision human and auditable.

The same fragmentation shows up on the patient side, and it is more visible there because the patient experiences it. Where patient-facing capability is built brand by brand, the same person exists as several disconnected identities across a portfolio, and continuity breaks at exactly the transitions that matter, when a product moves between formulations or indications and nothing underneath carries context across the boundary. Unifying every identity is not the answer either. Linkage has to be purpose-specific and has to enforce consent where applicable, contractual restrictions, retention, provenance and a hard boundary between patient support and promotion. HIPAA governs covered entities and business associates in covered contexts, and consumer health information outside those contexts can still fall under FTC or state requirements. A shared substrate is therefore a set of identity rules, purpose metadata, consent state, permitted linkages and auditable events. Sometimes those rules produce continuity. Sometimes they correctly produce separation.

## Build the foundation first, but build it through two workflows

Gartner predicts that by 2027 organizations prioritizing semantics in AI-ready data could improve agentic AI accuracy by up to 80 percent and cut operational costs by up to 60 percent. It separately predicts that more than 40 percent of agentic AI projects will be canceled by the end of 2027 on cost, unclear value or inadequate controls. Both are forecasts of what could happen. Neither is evidence from life sciences and neither proves the foundation-tax model. Read together they say something narrower and more useful: semantic context, demonstrated value and working controls have to arrive together or the project gets killed.

**This is the second place we think the industry has it backwards. The instinct under launch pressure is to fund the visible application and treat the foundation as overhead to be deferred.**

Built brand by brand, patient-context capability produces as many versions of the customer as there are brands, and every launch pays the foundation cost again. Built on shared substrate, each launch inherits what the last one established. The choice is not foundation or applications. It is sequence, and sequence decides whether launch capacity accumulates or resets to zero.

The practical route avoids both failure modes, the point-solution sprawl that adds one more disconnected tool per problem and the multi-year platform program that consumes capital before it changes a decision. Start with two bounded workflows that recur across launches. Data-request overlap detection and promotional content status visibility are good candidates. Stand up only the components those workflows need: agreed terms, source and rights metadata, approved-source references, identity and purpose rules, interfaces, access controls, lineage, decision ownership and event logging. Record the baseline before deployment. Keep human decision rights where they are. Judge speed and quality together, and reuse only the components that improve a second and a third workflow. NIST's voluntary AI Risk Management Framework offers a workable governance structure without dictating the architecture.

The foundation carries its own risks and they belong in the design. Sharing can broaden data use beyond permitted purposes. Hard-coded definitions freeze disagreements instead of resolving them. A shared component that fails takes more down with it. And chasing enterprise-wide completeness delays value indefinitely. Federated ownership, purpose-specific access, versioned definitions, local fallbacks and explicit rollback criteria are design decisions in their own right. So is adoption: in the 2025 Data and AI Leadership Executive Benchmark Survey, 92 percent of 125 executives named people and organizational change, ahead of technology, as the principal barrier to a data and AI driven culture. That is a perception measure from large organizations and proves nothing about pharmaceutical launch operations, but it matches what we see. Reusable systems survive only where owners trust the definitions and keep clear accountability.

## The claim

The industry knows what a good launch requires. What it has not solved is executing those requirements across overlapping portfolios without rebuilding the same data, content and control capability every time.

So here is the claim, and it is meant to be arguable. Within three years the biopharma companies that launch well will not be the ones with the best launch playbooks, because the playbooks are already shared and largely identical. They will be the ones whose second launch was cheaper than their first, and whose fifth was cheaper than their second, because each one left behind governed capability the next one inherited. Everything else is a launch team paying the foundation tax and calling it the cost of doing business.

## Disclosure and methodology

The authors are employees of Incedo Inc., which provides data, digital and AI services to life sciences organizations and may benefit commercially from the approaches discussed. Nitin Seth is Incedo's Co-Founder and Chief Executive Officer. Ashish Gupta is Executive Vice President, Global Head of AI & Data and Head of Life Sciences Practice. Anjali Gupta is Director, Life Sciences. Operational observations derive from Incedo client engagements in life sciences commercial data and digital organizations during 2025 and 2026, as reported by the authors. The number and distribution of engagements are not disclosed because of client confidentiality. The observations are qualitative and anonymized. They are not an independent prevalence study, a controlled evaluation or an audited outcomes analysis, and no client-identifying or patient-level information is presented. Any client results published later will report the unit of analysis, denominator, baseline, observation period, outcome definition and material exceptions.

# References


1. **Deloitte,** Measuring the Return from Pharmaceutical Innovation, 16th edition, May 2026.
2. **Health Advances**, Launch Costs: Spend Wisely, 2020. Fifteen single-product-company launches; commercialization spend derived from adjusted SG&A across a six-year window.
3. **McKinsey & Company,** The Secret of Successful Drug Launches. 210 new molecular entities, 2003 to 2009, against consensus forecasts available one year before launch.
4. **Deloitte Insights,** Key Factors for a Successful Drug Launch. 284 US launches, 2012 to 2021; 34 percent at or below 80 percent of expected launch-year sales; trajectory findings.
5. **US Food and Drug Administration,** CDER Novel Drug Approvals, annual reports 2021 to 2025; Office of Prescription Drug Promotion; 21 CFR 314.81(b)(3)(i).
6. **IQVIA,** The Rules of Loss of Exclusivity are Being Rewritten, July 2025.
7. **Veeva Systems,** Next-gen MLR Review with AI. Activities before formal review can add 5 to 150 days per asset.
8. **Indegene,** Is the MLR Review Process Slowing Pharma Marketing?
9. **Gartner,** press releases, 11 May 2026 and 25 June 2025, on semantics in AI-ready data and agentic AI project cancellation.
10. **Data & AI Leadership Exchange,** 2025 AI & Data Leadership Executive Benchmark Survey. 125 executives.
11. **US HHS,** De-identification of Protected Health Information; US FTC, Health Privacy business guidance.
12. **NIST,** AI Risk Management Framework.

# About the Authors

## Nitin Seth

Bestselling Author, Co-Founder & CEO

Nitin Seth is a global business leader, bestselling author, and a leading voice on AI-led enterprise transformation. His work spans the full arc of Digital, Data, and AI. As Co-Founder and CEO of Incedo Inc., he partners with Fortune 500 enterprises to reimagine their operating models and drive AI-native reinvention. His nearly three-decade career spans COO at Flipkart, Managing Director at Fidelity International, and Director at McKinsey. He is the author of the bestselling trilogy — Winning in the Digital Age, Mastering the Data Paradox, and Human Edge in the AI Age.

## Ashish Gupta

EVP, Global Head of AI & Data and Head of Life Sciences Practice

Ashish is the EVP playing dual role of Global Head of the AI & Data, and Head of Life Sciences Practice, where he partners with Fortune 200 clients to drive long-term value across the full spectrum of AI-led Business transformation. With 20 years of experience, he specializes in delivering innovative business solution across commercial excellence, patient engagement, payer access and clinical operations transformation across the Lifesciences value chain.

## Anjali Gupta

Director, Life Sciences

Anjali Gupta is Director of Life Sciences at Incedo, where she leads the firm's work with large biopharma clients on data, AI and commercial launch operations. She works with pharma data and digital organizations on the operating foundations behind product launches: data operations, content and MLR throughput, and customer intelligence. She holds a PhD and a Master's in Statistics.

**To learn more about how our products can drive your success, please email us at inquiries@incedoinc.com**

incedo

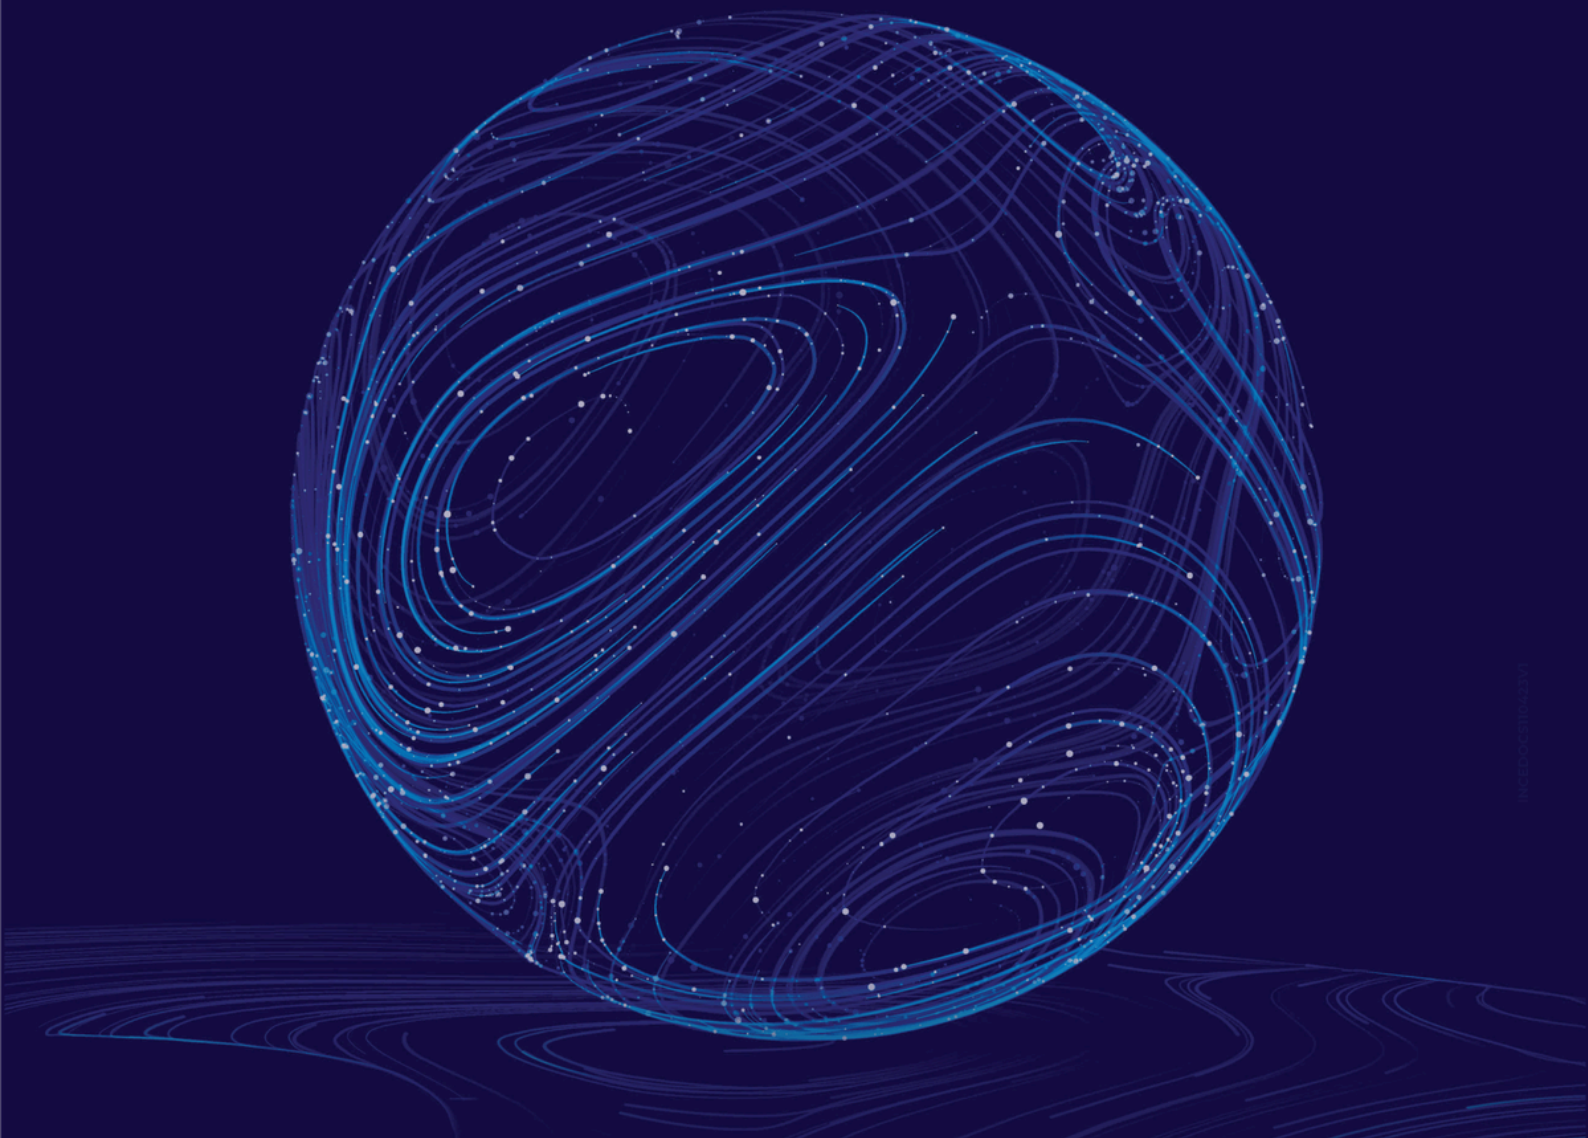

## About Incedo

Incedo is a global AI and data transformation specialist empowering companies to realize sustainable business impact from their digital investments by delivering ROI from AI@Scale.

As a long-term partner for strategy to execution, we operate at the intersection of business and technology. Our integrated services and platforms are built on the foundation of AI & Data, digital engineering, and operations transformation, bringing deep domain expertise and full stack capabilities together.

With over 4,000 people in the US, Canada, Latin America and India and a large, diverse portfolio of Fortune 500 enterprises and fast-growing clients worldwide, we work across banking & payments, wealth management, telecom, hi-tech and life sciences.

**10+**
**Fortune 500 Customers**

**10+**
**Global Locations**

**4k+**
**Employees Globally**

### Our Global Presence

**India**
Gurgaon
Chennai
Pune
Bengaluru
Hyderabad

**USA**
Santa Clara
New Jersey
Dallas
Boston

**Canada**
Ontario

**Mexico**
Guadalajara

**Follow us on:**

www.incedoinc.com