

Why good work goes unbuilt.

On pilots that never land, budgets that never hold, and the people, tools and places that never meet. Also: what changes when you write the test before you build.

The most expensive projects are the ones that never happened.

They also cost nothing, at least on paper. That is the trouble with them.

A grower knows where the food spoils. Two states away, an engineer has built a scheduler that would save it. Somewhere in between, a buyer would happily pay for the difference. (This is an illustrative example. You can probably supply a real one from your own line of work.)

The three never share a table. The idea is spoken, agreed with and left in the room. Next season the food spoils in the same place, right on time, which is at least one thing that runs on time.

Nobody writes this down as a loss, because nothing was built. There is no invoice for a tool that does not exist, and no line item for a conversation that went well and changed nothing. Accounting is very good at counting what happened. It is less good at counting what nearly did.

This paper is about that second column. It covers three ways good work goes unbuilt, and one fairly unglamorous fix: decide what would count as proof before anyone builds anything, then let the proof decide.

THREE WAYS IT GOES UNBUILT

- **Pilot purgatory.** The trial works. Nothing happens next.
- **The overrun.** The project happens, at a price nobody agreed to.
- **The missed meeting.** The people, the tools and the place never find each other.

WHO THIS IS FOR

Anyone who has watched a good idea get a standing ovation and then a filing cabinet. No engineering degree required, and no sales pitch until the last page.

The pilot worked. Nothing happened next.

~70%

of manufacturers were stuck in “pilot purgatory”: technology pilots that run for extended periods while viable technology never scales.

MCKINSEY, 2018 [1]

30%+

of generative AI projects predicted to be abandoned after proof of concept by the end of 2025.

GARTNER, 2024 [3]

78%

of organizations reported using AI in 2024, up from 55% the year before. Using is not the same as scaling.

STANFORD HAI, 2025 [4]

A pilot is supposed to be a question. Too often it becomes a place: somewhere a project goes to live, comfortably and indefinitely, with a dashboard. It is a strangely stable condition. Nobody has to decide to stop it, because nobody decided what success looked like in the first place.

McKinsey’s 2018 look at Industry 4.0 found that fewer than 30% of manufacturers were rolling those technologies out at scale.[1] The note worth underlining is the word *viable*. The technology in purgatory was not broken. It was simply never promoted.

Gartner’s reasons for abandoned AI proofs of concept read like a checklist of things that could have been agreed up front: poor data quality, inadequate risk controls, escalating costs, unclear business value.[3]

> *“Companies are drawn towards solutions that are exciting rather than solutions that are right.”*
>
> CHRISTOPH SCHMITZ, MCKINSEY, 2019 [2]

The same piece makes a second point that gets less airtime: advanced solutions often fail not because they produce wrong results, but because the people meant to use them do not understand or trust those results.[2]

That is a coordination problem wearing a technology costume. The fix is not a better pilot. It is agreeing, before the pilot starts, what result would earn the next step, and who gets to say so.

Over budget, over time, over and over again.

91.5%

of projects in Bent Flyvbjerg's database of 16,000+ go over budget, over schedule, or both.

FLYVBJERG & GARDNER, 2023 [6]

<1%

finish on time, on budget and deliver the benefits that were promised.

FLYVBJERG & GARDNER, 2023 [6]

79%

average cost overrun across 532 large projects, with delays averaging 52% against the initial timeline.

MCKINSEY, 2022 STUDY [7]

The headline above is not ours. It is what Bent Flyvbjerg, who has spent a career at Oxford collecting what happens to big projects, calls the "iron law of megaprojects."[5] In his 2014 overview he notes that nine out of ten megaprojects have cost overruns, and that the record has not improved over the seventy years for which comparable data exist.[5]

This is not a story about one kind of project or one part of the world. The database runs from Olympic Games to nuclear plants to IT systems, across 136 countries.[6] The habit, it turns out, travels well.

The interesting part is where the damage starts. McKinsey measured its overruns against estimates announced at the feasibility stage, and drew the plain conclusion: the trouble begins with poor cost and schedule estimates at approval.[7] The same article reports that collaborative contracts improved cost and schedule performance by 15 to 18 percent on average.[7]

Flyvbjerg and Gardner's prescription fits in four words: *think slow, act fast*.[6] Plan with rigor while plans are cheap to change. Build quickly once the plan holds.

Put another way, the most reliable route to finishing on budget is to do the arguing first, while arguing is still the cheapest thing on the schedule.

The people, the tools and the place rarely share a table.



Most unbuilt work is not blocked by a missing invention. The invention usually exists. It is in the wrong state, owned by the wrong company, or described in a vocabulary the people with the problem have never been taught.

Three groups have to meet: the people who live with a problem, the people who hold the knowledge and tools to address it, and a place where an answer can be tried without betting the farm, sometimes literally. Each tends to wait for the other two. Waiting is very cheap, until you add it up.

Shared platforms do not solve this on their own. In 2022 Maersk and IBM discontinued TradeLens, a shipping platform founded, in their words, "as an open and neutral industry platform."[8]

> *"While we successfully developed a viable platform, the need for full global industry collaboration has not been achieved."*
>
> MAERSK, ANNOUNCING THE END OF TRADELENS, 2022 [8]

The software worked. The table stayed half empty. The lesson is not that platforms fail. It is that neutrality has to be designed in, not announced, because competitors are reluctant to run their business through a rival's system, however good it is.

Line the three failures up and they share a shape. The test was never agreed, so the evidence could never decide.

Write the test first. Then build.

A clearing house finds the people. A lab finds the proof. That is the whole idea, and most of the work.

PrimusNeo is a clearing house plus an innovation lab. We match a real need with the knowledge and tools that can meet it, agree on how we will know if it worked, build a first version, and let the evidence choose what happens next. Our ethos is three words long, which is about the right length for an ethos:
Proof Over Promises.

01 **Understand the need** The people who live with the problem describe it, in their words.	02 **Form the team** Knowledge, tools and a delivery partner, each one named.	03 **Define the test** The question and its measures, set before anything is built.
04 **Build a first version** Existing tools where they fit. New ones where they don't.	05 **Review the evidence** Against the baseline, in the open.	06 **Put it to work** Expand it, adapt it, or stop. All three are allowed.

Step three is the one that matters, and the one most often skipped, because it is the step where people have to agree out loud. It turns "we'll know it when we see it" into a number, a baseline and a date.

The paradox is that defining the test in advance makes a project less certain to continue and more likely to deserve it. A pilot that is allowed to fail is the only kind that can properly succeed.

A project can end well in four ways: research people can use, a tool or service, a lasting partnership, or a new company. A clear, documented "no" is not on that list, but it is the cheapest outcome of all, and it frees the people and money for the next question.

> *"You cannot jump to the end. That is the point."*
>
> PRIMUSNEO HOUSE RULE

Many timelines. One line.

Coordination needs somewhere to live once the meeting ends. Under the lab sits a federated set of tools. A different organization stewards each part, and the parts meet through shared interfaces, file formats and permissions, so that no single company owns the table. In plain words:

OPQPWR PHYSICAL LAYER	Power, computing and storage in field racks a community can own, for places where the cloud is thin.	PARTNERSHIP STARTING
AMP OPERATIONAL RUNTIME	Sees and coordinates a physical site: maps, devices, who can go where. Works offline.	RUNNING AT TIE TOWN
NEOS AGREEMENTS & PERMISSIONS	Who decides what, written down so the rules outlive the people who wrote them.	COUNCILS AT TIE TOWN
AZOA IDENTITY, CONSENT, SETTLEMENT	Who agreed to what, and paying people for their work once. Not twice, and not never.	ALPHA
Fractal Engine ORGANIZATIONAL MEMORY	A record of who did what, kept on hardware you control. Plugins are refused unless they are allowed.	ALPHA
Elevate Impact RESEARCH ENABLEMENT	Our research partner. Turns a question into a study the evidence can answer, with the people affected at the table.	LIVE
OverSeasGenAI VENTURE SHIPYARD	For the rare result that deserves to become a company, judged on thresholds set before launch.	PARTNER

The workflow that joins these end to end (field data in, measured, decided, settled) is still in build. We would rather say so here than have you find it in a footnote. It is also in a footnote.[†]

Tie Town [TX]. A real place, which is the point.

One site in McDade, Texas. Not a case study with a thousand logos. One site.

Tie Town is a co-op and creator house in McDade, Texas, built to give new technologies, and the people making them, somewhere to work. It is the stack's first active operational deployment, which makes it our reference site: the place where each layer has to work on real ground before we claim it works anywhere else.

A reference site is a test you can visit. It turns "it should work" into "here is where it does, and here is where it doesn't yet." Ground truth with a receipt.

RUNNING THERE NOW

- AMP coordinating the site, including its media production.
- Programmatic access control for the creator facilities.
- Governance councils running on NEOS.
- Verifiable contributor memory through Fractal Engine (alpha).

WHAT WE ARE NOT CLAIMING

- That one site proves the model. It proves the parts can run on real ground.
- Outcome figures. We have not published any yet, so there are none here.
- That every layer is finished. Fractal Engine and AZOA are alpha.
- That the full connected workflow is live. It is in build.

LOCATION

30.28° N
97.24° W

Coordinates as listed in PrimusNeo's own materials.

WHY START WITH ONE SITE

Because the overrun research points the same way. Big things go better when they are built from small, repeatable pieces that have already been tested; Flyvbjerg and Gardner call it the Lego principle.[6] A town-sized test is one brick. It is less exciting than a grand launch, which, if you have read page 3, is rather the point.

Before you build anything, answer four questions.

1 **What would count as proof?**

Write it down: a measure, a baseline and a date. If it cannot be written down, it cannot be tested.

2 **Who gets to say it worked?**

And do they agree with the answer to question one? Find out now, not at the review.

3 **What does stopping look like?**

If stopping is not allowed, it is not a test. It is a ceremony with a budget.

4 **Who is missing from the table?**

The people with the problem, the people with the tools, or the place to try it. Usually one of the three.

If you can answer all four, you may not need us, which would be the best possible outcome for a paper like this. If you can't, or you have a problem worth solving, knowledge that could help, or a place to test a useful idea, we would like to hear about it.

GET IN TOUCH

primusneo.com

Sources & notes.

1 Enno de Boer, Diego Hernandez Diaz and Helena Leurent, “The Fourth Industrial Revolution and the factories of the future,” McKinsey Operations blog (originally published by the World Economic Forum), 17 August 2018. ~70% in pilot purgatory; fewer than 30% rolling out at scale.
mckinsey.com/capabilities/operations/our-insights/operations-blog/the-fourth-industrial-revolution-and-the-factories-of-the-future

2 Christoph Schmitz, “Hannover Messe 2019: The 3 P’s of Industry 4.0,” McKinsey Operations blog, 4 April 2019.
mckinsey.com/capabilities/operations/our-insights/operations-blog/hannover-messe-2019-the-3-ps-of-industry-40

3 Gartner, “Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025,” press release, 29 July 2024.
gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025

4 Stanford Institute for Human-Centered AI, *The 2025 AI Index Report*, April 2025 (organizational AI use drawn from McKinsey survey data).
hai.stanford.edu/ai-index/2025-ai-index-report

5 Bent Flyvbjerg, “What You Should Know About Megaprojects, and Why: An Overview,” *Project Management Journal* 45(2), 2014, pp. 6–19, doi:10.1002/pmj.21409. Author’s manuscript.
arxiv.org/abs/1409.0003

6 Daniel Akst, “Why do large projects go over budget?” *strategy+business*, 19 June 2023, reviewing Bent Flyvbjerg and Dan Gardner, *How Big Things Get Done* (Currency, 2023).
strategy-business.com/article/Why-do-large-projects-go-over-budget

7 Matthieu Dussud, Steffen Fuchs, Oliver Kernick, Piotr Pikul and Christoph Schmitz, “Seize the decade: Maximizing value through preconstruction excellence,” McKinsey, 19 September 2023 (reporting a 2022 McKinsey study of 532 projects of $100M+).
mckinsey.com/capabilities/operations/our-insights/seize-the-decade-maximizing-value-through-pre-construction-excellence

8 A.P. Moller - Maersk, “A.P. Moller - Maersk and IBM to discontinue TradeLens, a blockchain-enabled global trade platform,” 29 November 2022.
maersk.com/news/articles/2022/11/29/maersk-and-ibm-to-discontinue-tradelens

Notes. All sources were accessed on 9 November 2026. Figures are reported as each source states them; survey and database figures describe their own samples, not every project everywhere.

The grower, engineer and buyer on page 2 are an illustrative example, not a client. Descriptions of PrimusNeo, its stack and Tie Town [TX] come from PrimusNeo’s own materials (November 2026).

† **Status.** Live: Tie Town [TX]; Elevate Impact’s research module. Alpha: AZOA and Fractal Engine. OPQPWR is a new partnership; no on-premises deployment is running yet. The connected workflow between the layers is in build.