

Fizzle and the Practice-and-Platform Fix

Enterprise AI programmes are usually diagnosed at the model. The dataset was too small, the architecture was wrong, the metric was soft. On the engagement behind this whitepaper, a three-year subsurface-AI programme with a major operator in Oman, the model was never the thing that stalled. Data collection was. Midway through, with a downstream phase gated on a minimum of ten to 15 wells and wells arriving at a pace that put the timeline at risk, we stopped building and ran a data-management intervention, captured in a 39-slide workshop deck. What that deck argues, and what this whitepaper generalises, is that failed data initiatives share a pathology with a shape. It is a Situation-Complication-Resolution loop: a tactical fix for one team, a subject-matter-expert hire whose remit confirms the first fix, a data-strategy exercise that produces a strategy rather than access, a fall back to another tactical fix, and data debt that compounds every turn the loop runs. Underneath the loop sit three questions an executive can pose in any language and rarely gets answered: I cannot access the data, I do not trust the data in the report, and it takes too long to get to the data. We argue three things with three instruments. First, the fizzle is a loop, not a cliff, and its cost is the compounding debt, not the failed fix. Second, AI-readiness is not one switch but eight independent dimensions from the McKinsey MLOps frame, and a fizzle is what happens when a few are lifted and the rest are left at their before-state. Third, the exit is a practice-and-platform operating model delivered as three ordered four-day workshops, define requirements, design practices, architect the platform, sequenced across a three-horizon capability climb on a five-fiscal-year roadmap. The deliverable is a way for a board to recognise the pathology early and buy the fix that ends the loop, which is an operating model, not another model.

Tannistha Maiti

Senior AI Researcher

Tarry Singh

Founder & CEO

DATA-STRATEGY / ENTERPRISE-AI / OPERATING-MODEL / DATA-MANAGEMENT /
MLOPS-MATURITY / THREE-HORIZONS

CONTENTS

01	The failure has a shape, and the shape is a loop
02	Three questions the loop never answers
03	Why "AI-ready" is eight dimensions, not one switch
04	The vendor-ransom trap
05	The exit is a climb, and you enter it through workshops
06	What this is not: the pointer to the phased-operating-model piece
07	Reading the horizon against a real constraint
08	What a board should take from this
09	Limitations
10	References

A structural geologist can tell you within an afternoon whether a fracture-detection model is any good. Show them the picks, the dip, the azimuth, and they will find the interval where it lies to you. The model is legible, testable, and, when it is wrong, wrong in a way an expert can name. That legibility is a trap, because it draws every post-mortem toward the model and away from the place enterprise AI programmes actually stall, which is one layer down, where the data lives, or fails to.

We learned this on a three-year subsurface-AI programme with a major operator in Oman: a fractured carbonate play, image logs from two microresistivity imaging tools, a Detection Transformer that reached usable accuracy on fractures and beddings. The model worked. What did not work, for most of a year, was getting wells. A downstream phase of the programme was gated on a minimum of ten to 15 wells, and wells were arriving slowly enough that the timeline itself was at risk. So we stopped building models and did something a research team is not supposed to do in the middle of a research programme: we ran a data-management intervention, an interactive workshop, and put its argument into a 39-slide deck. This whitepaper is that argument, generalised past one carbonate reservoir to the question a board actually cares about: why do these initiatives fizzle, and what ends the fizzle?

The failure has a shape, and the shape is a loop

The word we kept coming back to was "fizzle," not "fail." A failure is an event. You can point at it, hold a review, assign a cause. A fizzle has no event. The programme does not crash; it loses pressure. Six months in, the dashboard nobody asked for is technically live, the specialist who was hired is technically productive, the data strategy is technically written, and the thing the programme was supposed to produce, decisions made on trusted data, is no closer than on day one.


That non-event has a structure, and naming the structure is the first useful thing a board can do. It runs as a Situation-Complication-Resolution loop, and the resolution is false, which is why the loop closes and runs again.

The **situation** is a tactical fix. A team has a data problem, so someone builds a point solution: one dashboard, one report, one pipeline for one question. It ships. It even helps, narrowly. Nobody is wrong to build it, and that is the problem, because a tactical fix that

helps is the most persuasive possible argument for building the next one instead of fixing the estate underneath.

The **complication** arrives in two moves. First, the organisation hires a subject-matter expert to own the data problem. The trouble is not the hire; it is the remit. The specialist is scoped to make the existing fix work better, so their expertise, honestly applied, confirms the tactical approach rather than questioning it. This is confirmation bias wearing an org chart. Second, someone senior notices the fixes are not adding up and commissions a data strategy. The strategy exercise produces a strategy. It produces a deck, a target-state diagram, a maturity model with the organisation plotted two rungs from the top. What it does not produce is access, because a strategy is a description of a destination, not a road to it.

The **false resolution** is the fall back. The strategy did not move data any closer, the specialist is still improving the original fix, and the pressure to show something reasserts itself, so the organisation falls back to what worked before: another tactical fix. The loop closes. And every turn of it leaves the same residue behind, which is data debt: one more point solution wired to one more silo, one more pipeline only one person understands, one more report whose numbers cannot be traced to their source.



The pathology of a failed data initiative, drawn as the loop it actually runs. A tactical fix (one dashboard for one team) invites a specialist hire whose remit confirms the first fix, which invites a data-strategy exercise that produces a deck rather than access, which invites a fall back to another tactical fix. Each turn of the loop leaves data debt behind, and the three questions the loop never answers stay open: I can't access the data, I don't trust the data in the report, it takes too long to get to the data. The teal funnel and debt bars are the loop; the plus and minus controls run it again and watch the debt index compound. The single orange block is the only element that argues: the exit is a practice-and-platform operating model delivered as three ordered workshops, define requirements, design data practices, architect the platform, not another model bolted onto the same broken estate. The four-stage loop, the three unfixed data problems, the three-workshop fix, and the Phase-3 floor of 10-15 wells are sourced from the engagement's mid-programme data-management workshop and the December board deck; the debt index and its per-turn compounding are an illustrative scale that dramatises the sourced fall-back beat, not a measured quantity.

The instrument above is the pathology drawn as the loop it actually is. The teal funnel is the four-stage cycle; the plus and minus controls run it again and watch the debt index compound. What the picture forces is a reframing of cost. The expensive thing about a fizzling initiative is not the tactical fix that did not scale. Tactical fixes are cheap, and some of them even earn their keep. The expensive thing is the compounding, because each unfixed turn raises the cost of the eventual real fix: more debt to unwind, more silos to reconcile, more one-person pipelines to reverse-engineer before anyone can lay a governed estate underneath. A board that sees the loop early pays to exit it once. A board that keeps funding the next tactical fix is paying, without naming it, to run the loop again.

Three questions the loop never answers

Underneath the loop sit three sentences. We heard versions of all three from executives who could not have told you a transformer from a decision tree, and their plainness is the point, because they are the actual test of whether a data programme is working.

The first is **I cannot access the data**. It lives in a system the person asking does not have credentials to, or in a format they cannot open, or behind an integration nobody finished. The second is **I do not trust the data in the report**. The number is on the slide, but its provenance is a shrug: no one can say which extract, which cut, which version of which table it came from, so a prudent executive discounts it. The third is **it takes too long to get to the data**. By the time the request has gone to the team that owns the pipeline, been queued, run, and come back, the decision it was meant to inform has been made on instinct.

Access, trust, time-to-data. Every tactical fix in the loop addresses one of these for one team for one question, and none of them addresses the three as a property of the organisation. That is the tell. A programme is fizzling precisely when it can point to a dozen places these questions were answered locally and no place they were answered generally. The workshop deck put it as a single blunt line, that data collection was moving at a really slow pace and the timeline would shift, and behind that operational line sat these three executive questions, unanswered at the level that mattered.


The three questions are also useful because they resist the flattery that maturity models invite. It is easy to feel good about a dashboard that answers "trust" for the sales team while the reservoir team still cannot open the file it needs. It is easy to celebrate a pipeline that cut one report from three days to three minutes while every other report still takes three days, because time-to-data is a distribution, not a headline, and a programme is only as fast as the request that blocks the decision actually waiting. On the engagement, the literal version of the access question was a well that would not download: the client's secure file transfer was slow enough that large log files restarted at 99% and had to be pulled again. That is not a glamorous failure, and no model architecture fixes it, but it was on the critical path, and a programme that ignores the unglamorous access failure in favour of the interesting modelling problem is choosing to fizzle.

THE BOARD-LEVEL TEST, IN THREE SENTENCES

Ask the three questions of the whole organisation, not one team: can a decision-maker reach the data, trust its provenance, and get it in time to decide? A programme that can only answer yes locally, one fix at a time, is fizzling, however busy it looks.

Why "AI-ready" is eight dimensions, not one switch

The reason the loop is so easy to stay inside is that "AI-ready" feels like a single state you either have or lack, and a tactical fix can always make it look, locally, like you have it. You do not. Readiness is a vector, not a switch, and the workshop borrowed the eight axes of that vector from the McKinsey frame for scaling AI: Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, and People and Skills. Each of the eight has a concrete before-state and a concrete after-state, and a fizzle is what you get when a few are lifted and the rest are left where they were.



AI-readiness is not one switch; it is eight independent dimensions, and an initiative that fizzles is usually one that lifted a few and left the rest at their before-state. The eight rows are the McKinsey Scaling AI like a tech native MLOps dimensions the engagement's data-management workshop adopted: Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, and People and Skills. Sweep the orange reading head down the axes with the arrows or by clicking a row; the selected dimension's before-state and the after-state the platform reaches are written in full on the right, with a teal maturity meter showing how far the after-state has moved. The eight dimensions and their before/after contrasts are sourced from the workshop and December board decks, including the deck's own worked example, live operations going from unstable solutions that go down for weeks at a time to monitored with instant alerts. The 0 to 100 after-state maturity index is an illustrative rendering of the sourced qualitative move, not a measured score.

Sweep the reading head down the eight axes and the shape of the problem shows. The deck's own worked example was Live Model Operations, before as "unstable solutions go down for weeks at a time" and after as "monitored with instant alerts." That gap is not a nuance; it is the difference between a system a business can lean on and one it cannot. The same before-to-after span runs across all eight. Data Management moves from data nobody can reach to a governed estate where the data a model needs is discoverable. AI Development moves from one-off notebooks to versioned datasets and pinned recipes, so a run is a record and not a memory. Technology Stack moves from whatever a vendor rents this year to a stack the operator can run and retrain on without a vendor. GRC and Security moves from governance-as-memo to signed monthly reports, RAID logs, and compliance

covering both wells data and personal data. People and Skills moves from capability that rents in and leaves with the consultant to in-country engineers who can run the platform after the engagement ends.

The load-bearing point is independence. These eight axes do not move together. A tactical fix can push AI Development a long way, a good dashboard can flatter Data Management, and a strong hire can lift People and Skills, all while Live Model Operations sits at "down for weeks" and GRC sits at "governance is a memo." The organisation feels more AI-ready and is not, because a chain is as strong as its weakest link and readiness is a chain. Any single axis left at its before-state is enough to keep the three executive questions unanswered, which is enough to keep the loop running. This is why we treat the maturity model as a checklist of eight independent gates rather than a single score to average. Averaging is exactly the move that lets a programme feel two rungs from the top while the rung that matters is on the floor.

The vendor-ransom trap

There is a specific way the loop hardens into something worse than a loop, and the deck named it without euphemism: organisations get "held ransom by expensive consultants and software vendors." The mechanism is not villainy. It is the natural consequence of answering the three questions with someone else's platform.

Each tactical fix that is bought rather than built takes a dependency. The dashboard is a vendor's dashboard, the pipeline runs on a vendor's stack, the integration is a consultant's integration, and the knowledge of how any of it works lives outside the organisation. Individually, each dependency is a reasonable buy-versus-build call. In aggregate, they compose into a position where the organisation cannot access, trust, or reach its own data without paying someone, and cannot stop paying without losing access to its own estate. That is the ransom: not a price, a lock. The Technology Stack axis, before as "whatever a vendor rents this year; the operator owns none of it," is where this trap is set, and it is why the after-state we hold out is a stack the operator can run and reason about on its own. On the Oman engagement this was not abstract. An independent, unbiased interpretation capability was worth building precisely because the alternative was renting it, at a cost, from service companies in perpetuity. Owning the capability was the point, and localisation, building the skills in-country so they stayed after we left, was how the ransom stayed unpaid.

The lock tightens because the dependencies compound the same way the data debt does. Each bought fix is not only a cost line; it is a piece of institutional knowledge that now lives outside the building. When the person who understood the vendor's integration leaves, the vendor is the only party who can change it, and the price of a change is whatever the vendor decides, because there is no second bidder for a system only the incumbent understands. Over a few years this converts a portfolio of reasonable buy decisions into a single supplier who holds the organisation's ability to reach its own data, and the switching cost is no longer a migration project but a capability the organisation never built and no longer knows how to build. That is why the escape is a People and Skills problem as much as a Technology Stack one: owning the stack is worth little if nobody in the building can run it.


None of this argues against ever buying. It argues that a programme run as a sequence of bought tactical fixes converges on the ransom by default, and that escaping it requires a deliberate posture toward ownership, set at the operating-model level, not decided one purchase order at a time.

The exit is a climb, and you enter it through workshops

If the fizzle is a loop, the fix cannot be a better point on the loop. It has to be an exit, and an exit is a different kind of thing: not another model, not another dashboard, but an operating model for the data capability itself. The workshop deck framed the exit as a three-horizon climb, borrowing the horizons from the growth literature and applying them to data rather than revenue.

Horizon 1 is defend the core: get access, get trust, cut time-to-data on the questions that already matter, so the organisation stops bleeding on the basics. Horizon 2 is build the emerging: the practices, the ownership, the RACI that make data a governed capability rather than a series of heroics. Horizon 3 is create options: a platform the operator owns and can extend into questions it has not asked yet. The horizons are not phases you finish and leave; they overlap, each maturing after the one before it, which is exactly the property that a tactical-fix organisation never achieves, because it keeps restarting Horizon 1 and never reaches Horizon 2.

You enter the horizons through three ordered workshops, each a structured four-day engagement with day-by-day activities, deliverables, and participants, laid on a five-fiscal-year roadmap from FY2022-23 to FY2026-27.



The resolution half of the argument, drawn as a climb. The fix for a fizzling data initiative is a three-horizon capability model applied to data, and the three horizons are entered by three ordered four-day workshops laid on a five-fiscal-year roadmap. Horizon 1 defends the core: Workshop 1 defines requirements and runs data discovery, fixing access, trust, and time-to-data on the questions that already matter. Horizon 2 builds the emerging capability: Workshop 2 designs data practices, RACI, ownership, reusable patterns. Horizon 3 creates options: Workshop 3 architects a data platform the operator owns. Drag the roadmap lever across FY2022-23 to FY2026-27 and the three horizon curves rise in sequence; the orange marker is the capability frontier, the envelope of what the operator can actually do at the swept fiscal year, and the left cards highlight whichever horizon is dominant then. The three-horizon model, the three four-day workshops, and the five-year roadmap are sourced from the engagement's data-management workshop and December board decks; the per-horizon S-curve heights are illustrative and render only the sourced peak-in-sequence ordering, not measured capability.

Drag the roadmap lever and the sequencing is the argument. Workshop 1 defines requirements and runs data discovery: the real questions the business needs answered and where the data to answer them actually lives, which is how you fix access, trust, and time-to-data on the things that matter rather than on whatever a single team shouted loudest about. Workshop 2 designs the data practices: the RACI that assigns ownership, the

patterns that make good data behaviour repeatable across teams instead of reinvented per project. Workshop 3 architects the data platform: the reference architecture that turns the practices into infrastructure the operator owns. The order is not decoration. Architecting a platform before you know the requirements is how you get an expensive answer to a question nobody asked, which is the strategy-for-strategy failure mode with a bigger budget. Designing practices before defining requirements is governance in a vacuum. The workshops are ordered because the failures they prevent are ordered.

There is one more thing the deck did that most vendor decks will not, and it belongs in any honest account of this fix. One architecture slide was greyed out and labelled as the area the project team was not aware of when the project commenced. It was an admission, on a client-facing slide, that the AI team had incomplete visibility into the client's own infrastructure until the heavy-engineering phase forced the issue. That admission is not a weakness of the method; it is the method working. A practice-and-platform operating model is the thing that surfaces the grey area early, in Workshop 1, rather than discovering it 18 months in when a model is ready and the data pipeline to feed it turns out to run through a system nobody on the team knew existed.

What this is not: the pointer to the phased-operating-model piece

A careful reader will notice that a practice-and-platform operating model and a phased capability build-out are close cousins, and they are. The mechanics of standing up an in-house subsurface-AI capability as a governed sequence of phases, with exit gates, staffing shapes, and a compute run-rate floor, are worked out in detail in our companion whitepaper, "A Phased Operating Model for Standing Up Subsurface AI Capability," and we will not re-derive them here. That piece answers the question "how do we build the capability, phase by phase." This piece answers a prior and more uncomfortable question: "why did the last three attempts fizzle before they got that far, and how does a board recognise the pathology in time to fund the fix instead of the next tactical patch." The operating model is the shared answer; the fizzle pathology is what makes the case for it.

Reading the horizon against a real constraint

The three-horizon climb can read as consultancy abstraction, so it is worth grounding it in the constraint that triggered the intervention in the first place. The programme needed ten to 15 wells to enter a downstream phase, and the wells were not coming fast enough. That is a Horizon 1 problem, precisely: access, in its most literal form. No amount of Horizon 3 platform architecture helps if the data does not arrive, and no data strategy deck makes a well appear.

The fix the deck recommended was not more technology. It was more of the right people on the client side, technical and domain resources dedicated to data collection and annotation, framed explicitly as also improving knowledge transfer and retention, which is the People and Skills axis and the Horizon 1 access problem answered together. This is what the practice-and-platform model looks like when it touches ground. The bottleneck was labelled data, and the fix was organisational: assign ownership, resource the collection, and treat the flow of trusted, labelled wells as the core to be defended before anything emerges on top of it. We can state the maturity relationship the climb encodes as a simple monotone requirement rather than a curve:

$$H_1 \preceq H_2 \preceq H_3, \quad \text{advance } H_{k+1} \text{ only once } H_k \text{ holds}$$

[Formula](#)[LaTeX](#)[Copy LaTeX](#)

Read the relation as ordering, not arithmetic: a later horizon may begin only once the one before it is holding, which is exactly the discipline a fizzling organisation lacks when it restarts Horizon 1 every quarter and never lets Horizon 2 mature. The instrument's S-curves render this ordering, each horizon peaking after the one before, and the frontier marker traces what the operator can actually do at a given point on the roadmap. The specific curve heights are illustrative; the ordering is the claim.

The five-fiscal-year span on the roadmap is worth dwelling on, because it is the single number most likely to make a board flinch. A data capability is not a quarter's work and it is not a project with a go-live date; it is a standing thing that matures over years, and the roadmap says so honestly. That length is not a reason to defer starting. It is a reason to start correctly, because the alternative to a five-year climb is not a faster climb, it is five years of the loop, which arrives at year five with more debt and no capability. The horizon

model reframes the timeline from a cost to be minimised into a sequence to be respected, and the practical consequence is that the first workshop can run in a week while the platform it points at takes years to mature, so the organisation feels progress on access and trust long before the full estate exists. Momentum on Horizon 1 is what buys the patience to let Horizon 3 arrive.

What a board should take from this

The value of naming the pathology is that it converts a vague unease into a decision. A board that has funded two or three data initiatives and watched each one lose pressure without failing outright is not unlucky. It is inside the loop, and the loop has a name and an exit.

The recognition test is the three questions, asked of the whole organisation and not one team: access, trust, time-to-data. If the honest answer is "yes, locally, in a dozen places, and no, generally, anywhere," the programme is fizzling regardless of how much activity it shows. The diagnostic is the eight-dimension readiness vector: not a single maturity score to feel good about, but eight independent gates, any one of which, left at its before-state, keeps the three questions unanswered. And the fix is an operating model, entered through ordered workshops on a three-horizon climb, that defends the core before it builds the emerging and builds the emerging before it creates options.

The one move that never works is the one the loop makes irresistible, which is to answer a fizzle with another tactical fix. It is cheaper this quarter and it deepens the debt for every quarter after. The practice-and-platform model costs more up front and ends the loop, and ending the loop is the only thing that turns three years of activity into a capability the organisation owns rather than a ransom it keeps paying.

WHAT THIS WHITEPAPER ARGUES

- 1 Enterprise AI initiatives usually fizzle at the data layer, not the model, and the fizzle is a repeatable Situation-Complication-Resolution loop: a tactical fix, an SME hire that confirms the fix, a data strategy that produces a deck instead of access, and a fall back to another tactical fix, with data debt compounding every turn.
- 2 The board-level diagnostic is three questions asked of the whole organisation, not one team: can a decision-maker access the data, trust its provenance, and get to it in time. A programme that answers yes only locally, one fix at a time, is fizzling however busy it looks.
- 3 AI-readiness is not one switch but eight independent McKinsey MLOps dimensions (Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, People and Skills); any single axis left at its before-state keeps the loop running, so the maturity model is a checklist of gates, not a score to average.
- 4 A programme run as a sequence of bought tactical fixes converges on a vendor-ransom lock, where the organisation cannot reach its own data without paying and cannot stop paying without losing access; escaping it is an ownership posture set at the operating-model level.
- 5 The exit is a practice-and-platform operating model, not another model, delivered as three ordered four-day workshops, define requirements, design data practices, architect the platform, sequenced across a three-horizon capability climb on a five-fiscal-year roadmap, grounded here in a real intervention on a major Oman carbonate engagement where the binding constraint was a ten-to-15-well data floor.

Limitations

This synthesis is drawn from a single intervention on one subsurface-AI programme with one operator, and its evidence is a mid-engagement workshop deck and a board deck, not a controlled study across many failed initiatives. The Situation-Complication-Resolution loop and the three executive questions are patterns we observed and found persuasive, not measured frequencies; a different sector or a different failure mode may not fit the shape. The eight readiness dimensions are adopted from a published McKinsey frame and inherit whatever is contestable in that frame, and the before-and-after contrasts are qualitative

characterisations of the engagement's own cells rather than scored assessments. The three-horizon model is likewise a borrowed lens applied to data capability; the instrument's per-horizon curves and the debt index in the pathology instrument are illustrative devices that render sourced ordering and sourced beats, not measured capability or measured debt. Finally, the workshop-driven operating model is presented as the fix we ran and recommended on this engagement, not as a validated intervention with a measured success rate; its case rests on the coherence of the diagnosis and the honesty of the account, including the greyed-out slide where the team admitted what it did not know, rather than on a controlled comparison against the tactical-fix path.

References

McKinsey & Company, 2021. Scaling AI like a tech native: The CEO's role. The eight-dimension MLOps maturity frame this whitepaper adopts, spanning data management, AI development and deployment, live model operations, technology stack, GRC and security, ML assets management, and people and skills.

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/scaling-ai-like-a-tech-native-the-ceos-role>

Baghai, M., Coley, S., White, D., 1999. The Alchemy of Growth: Practical Insights for Building the Enduring Enterprise. The three-horizon growth model this whitepaper applies to data capability, defend the core, build the emerging, create options.

<https://www.hachettebookgroup.com/titles/mehrdad-baghai/the-alchemy-of-growth/9780738203027/>

Access, Trust, Time-to-Data: Why Enterprise AI Initiatives Fizzle and the Practice-and-Platform Fix

Published August 2026. Generated 2026-09-08.

<https://www.earthscan.io/whitepapers/access-trust-time-to-data-why-enterprise-ai-initiatives-fizzle>

Copyright EarthScan. All rights reserved.