EARTHSCAN

WHITEPAPER 18 PAGES / AUGUST 2026

The Roadmap Behind the Models: Pitching an Operator a Data Operating Model, an In-House Academy, and a Joint Venture

Most of the record of this engagement is about the machine-learning pipeline. This whitepaper is about the other half of the 2022-2023 proposal round: an 11-domain data-management roadmap bucketed across three horizons, a per-phase KPI and business-impact contract, an eight-dimension organizational maturity diagnosis, and two board constructs that most vendors never table at all. An in-house data-science academy for the operator's own staff, and joint-venture planning. Read together, the pitch treated AI delivery as capability transfer, not model delivery, and the roadmap board is the argument for why that framing changes what an operator is actually buying.

Tarry Singh
Founder & CEO

Quamer Nasim
ML Research Engineer

DATA-OPERATING-MODEL / DAMA / DATA-MANAGEMENT-ROADMAP / CAPABILITY-TRANSFER / DATA-GOVERNANCE / DATA-SCIENCE-ACADEMY

CONTENTS

The visible artefact of an applied-AI engagement is a model. A fracture detector, a vug quantifier, a well-to-well correlation tool. Those are the things a board can point at, the things that end up in a conference paper, the things a competitor can benchmark. They are also, in a two-year engagement with a major operator in Oman, roughly half of what was actually pitched.

The other half almost never makes the write-up, because it does not photograph well. It is a spreadsheet of data-management domains bucketed into time horizons, a table pairing each project phase with a measurable KPI and the business impact it was supposed to buy, an eight-dimension diagnosis of how the operator's own organization handled data before and after the work, and two line items that most AI vendors would never put in front of a client at all. A proposal to stand up an in-house data-science academy inside the operator, and a line reading, plainly, Joint Venture Planning.

This whitepaper is about that other half. The claim it makes is narrow and, we think, load-bearing for anyone commissioning subsurface AI. The advisory pitch that sat alongside the machine-learning pipeline treated AI delivery as capability transfer, not model delivery. It was designed so that at the end of the engagement the operator would own a data operating model, a trained bench of its own people, and a governance cadence. Not just a set of trained weights it would have to keep paying someone else to retrain. Two of our earlier pieces cover the downstream halves of that story: the ICT transition that left the operator running its own compute, and the phased engagement operating model that sequenced the whole thing. This paper is about the artefact that came first, the roadmap that made both of those downstream steps legible to a board before a single well of data had moved.

The spreadsheet nobody photographs

The roadmap lived in an ordinary Info-Tech roadmap tool, an Excel template with a bubble-chart output and three time buckets: short-term at 6 months, medium-term at 6 to 12 months, long-term at 12+ months, with room for up to 26 items in each horizon. What went into it is what matters. Not machine-learning tasks. Data-management domains, eleven of them, drawn straight from the DAMA-style vocabulary that a data-management professional would recognize on sight.

The eleven: Data Asset Management Strategy. Data Platform Selection and Implementation. Data Quality Management. MetaData Management. Master Data Management. Data Warehouse Management. Data Integration Management. Data Governance Management. Data Architecture Management. Develop AI Technology Standards. Develop Staff Skills.

Read that list against what the engagement was nominally about, which was picking fractures and beddings off borehole image logs, and the gap is the whole point. None of those eleven items is a model. Nine of them are the plumbing an organization needs before a model is worth training, and the tenth and eleventh (AI technology standards and staff skills) are about the operator's own capacity to keep doing this after the vendor leaves. The roadmap tool tags each item with a strategic-goal category. Four were in use: Operational Excellence, New Services, Digital Transformation, and a category simply named Joint Venture. It also carried a status vocabulary a program manager would want, budget states (already budgeted, pending approval, planning) and approval levels (executive, director, manager), so the board reading the bubble chart could see not just what was proposed but where each item sat in its own approval pipeline.

Two items on that roadmap are not data-management domains at all, and they are the reason this paper exists. One is an in-house DataScience Academy for the operator. The other is Joint Venture Planning. Put an academy and a JV on the same board as data-quality management and metadata management, and you have stopped pitching a project. You are pitching an operating model, and a relationship.

Reading the roadmap as an argument

The roadmap board below plots all of it. Eleven data-management domains in teal, bucketed across the three horizons, each carrying its strategic-goal category. And the two board constructs (the academy and the joint-venture line) drawn as the single orange thread, because they are what turns the whole thing from a delivery plan into a capability-transfer plan.

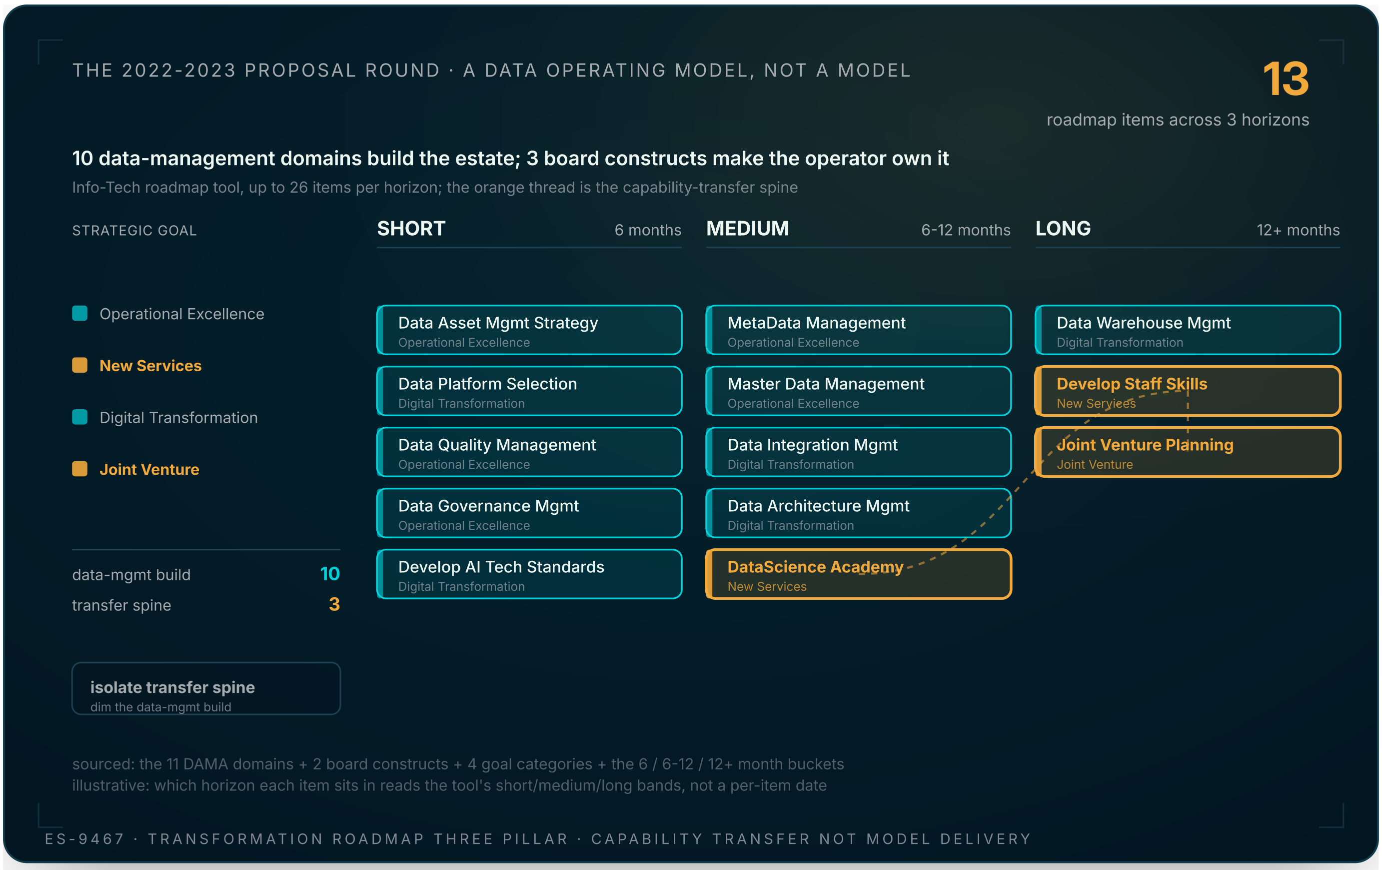

The 2022-2023 proposal round read as capability transfer, not model delivery. Eleven DAMA-style data-management domains (from data-asset strategy and governance through metadata, master data, integration, architecture and staff skills) are bucketed across an Info-Tech roadmap tool's three horizons: short (6 months), medium (6-12 months), long (12+ months), up to 26 items per horizon. The teal pills are the data-estate build; the single orange thread is the capability-transfer spine that turns the whole roadmap from a handed-over model into an operator-owned capability: an in-house DataScience Academy, staff-skills development, and Joint Venture Planning. Each pill carries its strategic-goal category (Operational Excellence, New Services, Digital Transformation, Joint Venture). The toggle isolates the spine so the argument is unmissable. The 11 domains, the 2 board constructs, the 4 goal categories, and the three horizon-bucket definitions are sourced from the engagement's roadmap tool and proposal round; the horizon in which each individual item sits is an illustrative reading of the tool's short/medium/long bands, not a per-item sourced date.

The design choice in that instrument is deliberate, and it is the same choice the original roadmap made. The teal items are necessary and unglamorous. You cannot run models on data you cannot trust, and Data Quality Management, Data Governance, Master Data, and Metadata are the domains that make data trustworthy. They sit early, in the 6-month and 6-to-12-month horizons, because they gate everything after them. But teal alone is a data-cleanup project, and a data-cleanup project is something an operator buys once and forgets. The orange thread is what makes the roadmap about the operator rather than about us. The academy trains the operator's own staff. The staff-skills item develops them further. Joint Venture Planning contemplates a structure in which the operator is a co-owner of the capability, not a customer of it. Those three sit later, in the 6-to-12 and 12+ horizons, because you cannot transfer a capability that does not yet exist. First you build the estate, then you hand over the keys.

That ordering is the argument. Capability transfer is not a slogan you add at the end of a statement of work. It is a sequence, and the sequence has to be visible on a board's single sheet of paper or it does not get funded. The horizon placement of each item on our board is an illustrative reading of the tool's short, medium, and long bands rather than a per-item sourced date, but the buckets, the eleven domains, the two constructs, and the four categories are exactly what the roadmap carried.

Why an operator would want an academy on the roadmap

The instinct in most AI engagements is to keep the client dependent. The vendor's revenue model rewards it. If the operator cannot retrain the model on next year's wells without you, you have an annuity. Putting a DataScience Academy on the roadmap is a decision to give that annuity up on purpose, and it is worth being honest about why it made sense here rather than pretending it was pure altruism.

The operator was in Oman, and Omanization (the national policy of building local employment and skills into contracts) is not a nice-to-have there. It is a procurement reality. A vendor that shows up with a plan to train the operator's own nationals, in partnership with a local academic institution, is speaking the language the operator's own board answers to. The academy line was not charity. It was the item that made the rest of the roadmap fundable, because it aligned the AI program with the operator's obligations to develop its own people. That the academic partnership ran through a local university, and that the year-two engagement later formalized six live training sessions (three online, three on the operator's own premises) with an AI-and-geology curriculum, is the downstream evidence that the academy line was real and not decorative.

There is a second reason, less about policy and more about physics. A model trained on 14 wells is a model that will drift the moment the fifteenth well comes in with a slightly different picking style, as we learned the hard way when adding a fifteenth well actually dropped fracture F1 at a 5 cm tolerance from 60% to 57% on the validation set. More data hurt, because the label style differed. A capability that can diagnose and fix that kind of regression is not a set of weights. It is a trained person who understands why it happened. The academy is the only line on the roadmap that produces that person.

The eight dimensions the workshop actually measured

The roadmap did not arrive in a vacuum. It came out of a diagnosis, run as an interactive workshop at the point in the engagement where data collection had stalled and timelines were visibly slipping. The blunt warning in that room was that Phase 3 needed a minimum of 10 to 15 wells to be worth starting, and the wells were not arriving fast enough. The workshop's job was to explain why, and the explanation was not about the models. It was about the organization.

The diagnostic instrument was an eight-dimension maturity framework, borrowed from McKinsey's "Scaling AI like a tech native: The CEO's role," and it measured the operator's own MLOps maturity before and after the intended work: Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, and People and Skills. Each dimension came with a concrete before-and-after contrast rather than a score. The before state for live operations was memorable enough to quote from the deck: unstable solutions that go down for weeks at a time. The after state was monitored operation with instant alerts.

The eight-dimension before/after MLOps maturity framework the advisory ran as a workshop when data collection had stalled, sourced from McKinsey's 'Scaling AI like a tech native: The CEO's role'. Each dimension (Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, People and Skills) is diagnosed at a current state and a target state; the toggle steps every dimension from before to after. The point is that the intervention assessed the operator's organization, not the models, which is what capability transfer means in practice. The single orange axis is People and Skills, the dimension the in-house DataScience Academy and the staff-skills roadmap items were built to move. The eight dimensions and their before/after framing are sourced from the workshop deck; the 1-to-5 maturity levels are an illustrative encoding of the deck's qualitative before/after contrasts, not sourced scores.

The lens above steps each of those eight dimensions from its current state to its target state. The reason this instrument belongs in a paper about capability transfer is that seven of the eight axes are organizational, not algorithmic. You do not move Data Management, GRC and Security, or ML Assets Management by training a better network. You move them by changing how the organization works. And the eighth axis, People and Skills, is the one the academy and the staff-skills roadmap items were built specifically to move. We have drawn it orange for that reason. The numeric 1-to-5 levels on the lens are an illustrative encoding of the deck's qualitative before-and-after contrasts, not sourced scores, but the eight dimensions and their framing are exactly what the workshop put on the wall.

The workshop also named the three failure modes it was trying to prevent, and they are worth stating because they are the failure modes of every stalled data program. First, the tactical fix: someone patches the immediate problem, which relieves the pressure and removes the reason to fix the underlying cause. Second, the confirmation loop: the organization hires a subject-matter expert, the expert confirms the organization's existing beliefs, and no value is created. Third, the strategy-for-its-own-sake spiral, where the organization writes a data strategy, fails to execute it, and falls back to tactical patching, accumulating data debt each cycle. The roadmap was the antidote to all three, because a roadmap with horizons and approval states is the thing that survives the temptation to patch and move on.

The three complaints the roadmap was built to answer

Before any of the eleven domains made sense, the workshop had to establish that the operator had a data problem at all, and it did so with three sentences that any large organization will recognize. The deck framed the operator's situation as three key data problems, phrased from the point of view of the people inside the business: I cannot access the data. I do not trust the data in the report. It takes too long to get to the data. Those three complaints are the entire justification for a data-management roadmap, and they are worth dwelling on because each one maps to specific domains on the board.

I cannot access the data maps to Data Integration Management and Data Architecture Management, the domains that make data reachable across a fragmented estate. I do not trust the data in the report maps to Data Quality Management, Master Data Management, and Data Governance, the domains that make data believable. And it takes too long maps to Data Platform Selection and Data Warehouse Management, the domains that make data fast. Stated that way, the roadmap is not a menu of best practices imported from a framework. It is a direct response to three things the operator's own people were already feeling, which is why it could be funded. A roadmap that answers a felt problem gets a budget line; a roadmap that answers a framework gets a polite thank-you.

The workshop was also candid about a failure mode specific to how organizations buy this kind of work. It named, in as many words, the risk of being held ransom by expensive consultants and software vendors, and it did so in a deck that was itself a consultant's pitch. That is an unusual thing to put on a slide. It is also the tell that the framing was capability transfer rather than lock-in, because a vendor planning to become the expensive consultant does not warn the client about expensive consultants. The academy, the staff-skills stream, and the self-service inferencing KPI are the concrete answers to that warning: each one is a mechanism by which the operator becomes less dependent on outside help over time, not more.

There was an honest admission buried in the same diagnosis that is worth surfacing, because it cuts against any temptation to read the roadmap as a vendor's tidy self-portrait. One architecture slide greyed out a section of the operator's own data estate and labelled it, plainly, an area the project team was not aware of when the project commenced. That is a vendor admitting, on the record, that it did not fully understand the client's infrastructure until the heavy-engineering phase forced the issue. It belongs in a paper about capability transfer precisely because it is the opposite of a lock-in posture. A vendor optimizing for dependency hides its blind spots. A vendor optimizing for transfer draws a grey box around them and says, here is what we did not know, so that the operator's own team knows where to look.

Contracting a capability, not a deliverable

A roadmap and a maturity diagnosis are advisory. What made the capability-transfer framing binding rather than aspirational was the KPI and business-impact contract that sat under each phase. This is where the pitch stopped being a slide and became a set of measurable commitments the operator could hold us to.

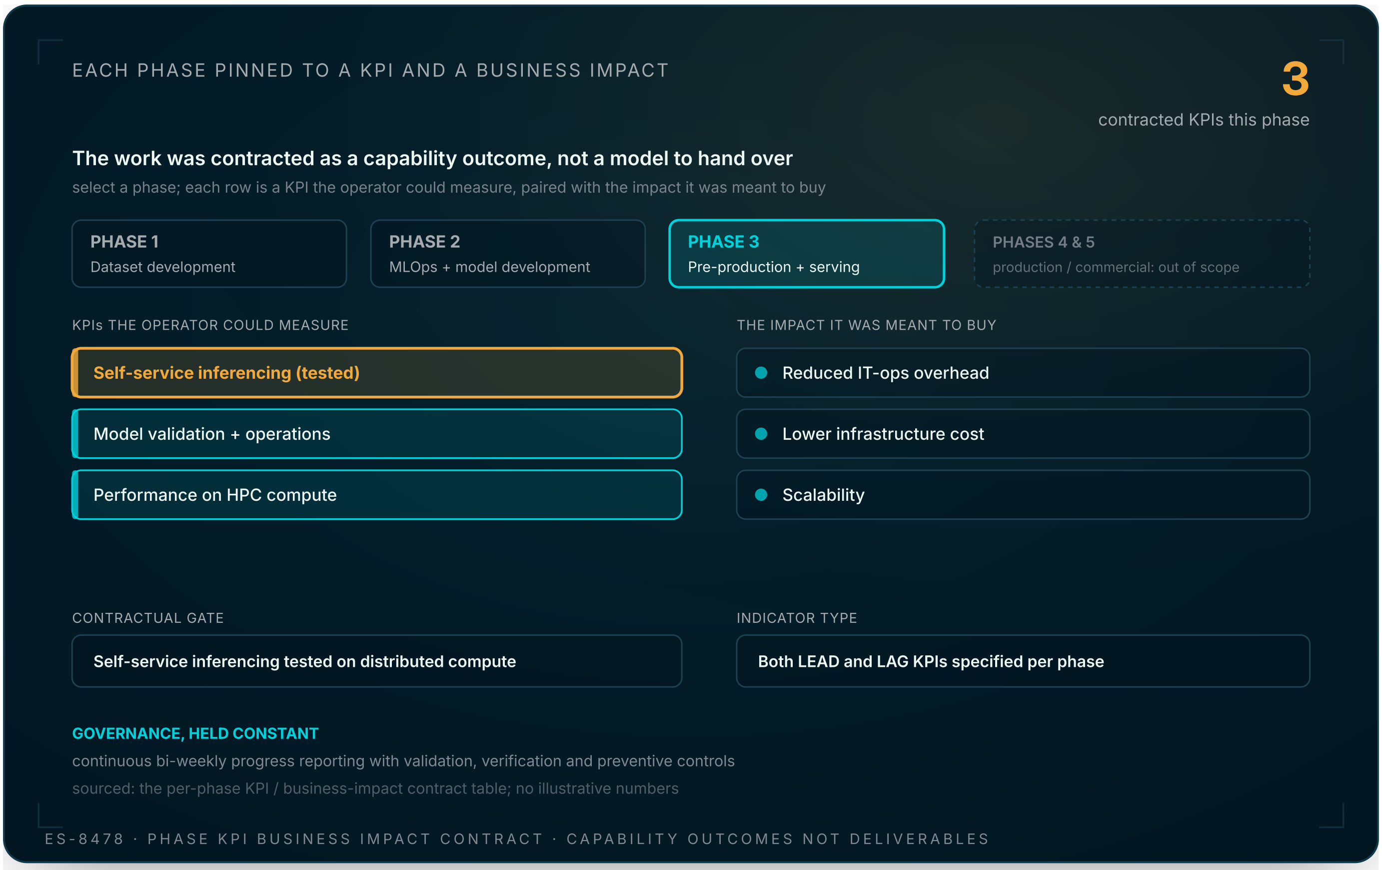

The per-phase KPI and business-impact contract table that sat under the engagement, read as evidence that the work was contracted as a capability outcome rather than a model to hand over. Select a phase to see its KPIs (the things the operator could measure) paired with the business impact each was meant to buy, plus the phase's contractual gate. Phase 1 pins Time to Deploy and dataset Standardization against reduced Time to Market; Phase 2 pins new algorithms, models, versioning and performance against rapid iteration and pipelining, gated at at least one functioning algorithm and one functioning model; Phase 3 pins self-service inferencing, model validation and operations against reduced IT-ops overhead and lower infrastructure cost. The single orange KPI is Phase 3 self-service inferencing, the lead indicator the contract used for the operator's own production footprint, the point at which the roadmap becomes about the operator running its own AI. Governance is held constant underneath: continuous bi-weekly reporting with validation, verification and preventive controls. Phases 4 and 5 (production and commercial) were explicitly out of scope. Every item is sourced from the contract table; there are no illustrative numbers.

The contract table above pairs each phase's KPIs (the things the operator could measure) with the business impact each was meant to buy. Phase 1, dataset development, was gated on Time to Deploy within the phase and on Standardization, defined as a complete, ready dataset for the selected wells, with the stated impact of reducing Time to Market. Phase 2, the MLOps and model-development phase, was gated on new algorithms, new models, model versioning, and model performance, with a hard contractual floor of at least one fully functioning algorithm and at least one fully functioning model with appropriate versioning, and the impact framed as rapid iteration, model optimization, and model pipelining. Phase 3, pre-production and serving, was gated on self-service inferencing (tested), model validation, model operations, and performance on distributed HPC compute, with the impact framed as reduced IT-ops overhead, lower infrastructure cost, and scalability.

The single most telling KPI is the Phase 3 one: self-service inferencing, tested. The contract labels it a lead indicator for the operator's own operational and cost footprint in production. A lead indicator, not a lag one. That is a small piece of contract vocabulary carrying a large idea. The engagement was not measuring whether we delivered a model. It was measuring whether the operator could run inference on its own, without us in the loop. A model-delivery contract has no such KPI, because in a model-delivery contract the vendor stays in the loop by design. The whole table is drawn from the per-phase KPI and business-impact document; there are no illustrative numbers in it.

Two structural details reinforce the framing. The whole thing ran on a governance cadence of continuous bi-weekly progress reporting with validation, verification, and preventive controls, which is the reporting rhythm of an organization being taught to govern its own program, not of a vendor filing status updates. And Phases 4 and 5, production and commercialization, were explicitly held out of scope, deferred pending the final presentation. That boundary matters: the engagement scoped itself to building the capability and stopped at the point where the operator would take it into production on its own terms. It did not try to own the production phase. Which is, again, what capability transfer looks like when it is contracted honestly rather than promised loosely.

The joint venture, and what it says about the ambition

The last item on the roadmap is the one that reveals the true altitude of the pitch. Joint Venture Planning, tagged with its own strategic-goal category. Not a service line. Not a support contract. A construct in which the operator and the AI partner would share ownership of the capability being built.

We should be careful here, because the JV line was planning, not a signed deal, and the anonymized record does not support claiming an executed structure. What it does support is the observation that the JV line existed on the same board as data-quality management, which tells you how the engagement was being framed at the top. A vendor pitching a project does not put a joint venture on the roadmap. A partner contemplating a shared, ongoing capability does. The academy trains the people, the roadmap builds the estate, the

KPI contract makes the capability measurable, and the JV line is the vehicle by which the resulting capability could be co-owned rather than sold. Read the four together and the model-delivery reading of the engagement collapses. You cannot joint-venture a set of weights. You can only joint-venture a capability that keeps producing value, which is exactly what the other three artefacts were built to create.

There is a commercial honesty worth naming in the JV framing too. Some of what the roadmap contemplated (an independent, unbiased dip dataset, and a vug-dimension methodology) was flagged in the same period as commercially novel, the kind of thing service companies charge for or that was, per the internal framing, not available anywhere. A JV is the natural structure for jointly owning an asset like that, rather than one party selling it to the other. The JV line and the commercial-novelty framing are two views of the same underlying bet: that the capability being built was worth co-owning.

Reporting built for three different rooms

One more artefact from the same period makes the capability-transfer intent concrete in a way that is easy to overlook, because it is about slide decks rather than models. The board reporting was deliberately structured for three distinct audiences, sectioned that way in the deck itself: the domain experts (the geoscientists doing geology and geophysics), the CIO, CDO, and IT function, and the CEO, managing director, and board. Same program, three lenses, because the three rooms need different things to be true.

This matters more than it looks. A vendor that reports one way, in its own idiom, is a vendor teaching the client to consume the vendor's output. A vendor that reports in three idioms is a vendor teaching three different parts of the client organization to own three different parts of the program. The geoscientists needed to see sinusoid-count parity against their own picks and the intervals where the model flagged a feature they had not. The CIO and CDO needed to see the data-management maturity movement and the infrastructure tiering. The board needed to see the growth-horizon story and the strategic-goal categories that tie back to the roadmap. Building all three, every reporting cycle, is more work than building one. It is the work capability transfer requires, because a capability the board does not understand does not get renewed, and a capability the geoscientists do not trust does not get used.

The growth-horizon story told to the board is worth naming for what it reveals about altitude. The deck used a three-horizon growth model (defend the core, build the emerging, create the options) to place the AI program not as a project but as an option-creating bet, and it paired that with a culture-versus-financial-performance argument to make the case that the organizational change was the point. Those are board-language constructs. They are not how you sell a fracture detector. They are how you sell an operating model, and the fact that they sat in the same reporting architecture as sinusoid-count tables is the clearest evidence that the engagement was being run as a transfer of capability across the whole organization, geologists and executives alike.

The governance scaffolding under all of it was equally deliberate: monthly signed reports, daily stand-ups, weekly meetings, and RAID logs tracking risks, actions, issues, and decisions, with GDPR compliance covering both personal data and the wells data itself. RAID logs and signed monthly reports are not how a vendor documents its own work for its own files. They are how an organization is taught to govern a program it intends to keep running. Every one of those cadences is a habit being installed in the operator, not a deliverable being produced for it.

What an operator should take from the roadmap board

The reason this artefact deserves a whitepaper of its own, rather than a footnote in the model papers, is that it inverts the default question an operator asks when commissioning AI. The default question is "will the model work?" The roadmap reframes it to "will we be able to run this ourselves in two years?" Those are different questions, and they lead to different procurement decisions.

For a chief data officer or a CIO evaluating a subsurface-AI pitch, the roadmap board gives a concrete checklist, and it is the checklist we would apply to our own pitches:

- Does the proposal put data-management domains (governance, quality, metadata, master data, architecture) on the roadmap, or only models? A pitch that is all models is a pitch that assumes your data is ready, and it almost never is.
- Is there a line item that develops your own people, and is it early enough on the roadmap to matter? An academy or a staff-skills stream in the 12+ horizon is decoration; one that starts building the bench while the estate is being built is real.

- Is each phase pinned to a KPI you can measure and a business impact you can point at, with at least one lead indicator for your own operational footprint? Self-service inferencing, tested, is the tell. If the contract has no KPI that measures your independence, the vendor is not planning to leave.
- Does the governance cadence teach your organization to run the program, or just report to you? Bi-weekly validation and verification is a teaching rhythm; monthly status decks are not.
- And, at the top of the sheet, is the vendor contemplating a shared future (a JV, a co-owned asset) or a support annuity? The answer tells you what they think they are selling.

None of these questions is about accuracy. Every model paper we have written on this engagement argues its accuracy carefully, and accuracy matters. But an operator that optimizes only for the model's F1 and ignores the roadmap ends up with a good model it cannot maintain, which within a year or two is a bad model. The roadmap board is the artefact that argues for the other thing: that what an operator is really buying, when it buys subsurface AI well, is a data operating model, a trained bench, and a governance habit, with the model as the visible proof that the capability works.

WHAT THIS WHITEPAPER ARGUES

1. The 2022-2023 proposal round pitched far more than models: an 11-domain DAMA-style data-management roadmap (governance, quality, metadata, master data, integration, architecture, warehouse, platform, asset strategy, AI standards, staff skills) bucketed across three horizons on an Info-Tech roadmap tool (short 6 months, medium 6-12 months, long 12+ months, up to 26 items per horizon).
2. Two board constructs turn the roadmap from model delivery into capability transfer: an in-house DataScience Academy for the operator's own staff, and Joint Venture Planning. They sit in the later horizons because you build the data estate first, then hand over the keys.
3. The academy line was fundable because it aligned the AI program with Omanization and a local academic partnership, and because a model trained on 14 wells drifts (adding a 15th well dropped fracture F1 at 5 cm from 60% to 57% on validation), so the durable asset is a trained person, not a set of weights.
4. An eight-dimension MLOps maturity workshop (McKinsey 'Scaling AI like a tech native') diagnosed the operator's organization, not the models, across Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, and People and Skills, run when data collection had stalled and Phase 3 needed a minimum of 10-15 wells.
5. The capability-transfer framing was contractual, not aspirational: a per-phase KPI and business-impact table pinned Phase 1 to Time to Deploy and Standardization, Phase 2 to a floor of >=1 functioning algorithm and model, and Phase 3 to self-service inferencing (tested), which the contract labels a lead indicator for the operator's own production footprint, all governed by continuous bi-weekly reporting with validation, verification and preventive controls.

References

McKinsey & Company, 2020. *Scaling AI like a tech native: The CEO's role.* The source of the eight-dimension MLOps maturity framework (Data Management, AI Development, AI Deployment, Live Model Operations, Technology Stack, GRC and Security, ML Assets Management, People and Skills) that the mid-engagement workshop used to diagnose the operator's organization before and after the intended work. https://www.mckinsey.com/

Limitations

This paper reads a proposal-round artefact, and proposal-round artefacts describe intent, not outcomes. The roadmap tool records what was pitched across three horizons; it does not by itself prove that every one of the eleven data-management domains was executed on the horizon shown, and the horizon placement of individual items in the accompanying instrument is an illustrative reading of the tool's short, medium, and long bands rather than a per-item sourced date. The Joint Venture Planning line was planning, and the anonymized record does not support any claim that a joint venture was executed; the value of the line is what its presence says about how the engagement was framed, not evidence of a signed deal. The eight-dimension maturity levels shown in the lens are an illustrative 1-to-5 encoding of the workshop deck's qualitative before-and-after contrasts, not sourced numeric scores; only the dimensions themselves and their framing are sourced. The KPI and business-impact table, by contrast, is reported directly from the per-phase contract document and carries no illustrative numbers. Client identity and specific well identifiers are withheld throughout; the operator is described only as a major operator in Oman, and the academic partner and the tools involved are named where the engagement record and current disclosure policy permit. Finally, this is one engagement. The checklist we draw from it is a considered generalization, not a benchmarked finding, and an operator applying it should weigh it against its own procurement context.

EARTHSCAN

The Roadmap Behind the Models: Pitching an Operator a Data Operating Model, an In-House Academy, and a Joint Venture

Published August 2026. Generated 2026-09-08.
https://www.earthscan.io/whitepapers/the-roadmap-behind-the-models-data-operating-model-academy-joint-venture