

After the Model Ships: Anatomy of a Year-Two AI Technology-Transfer Engagement

Most write-ups of an AI engagement stop when the model ships. The more revealing document is the one that comes next: the services agreement for the year after, which decides whether the operator ends up owning a capability or renting a dependency. We read the year-two agreement from a subsurface-AI programme with a mid-sized Middle East carbonate operator as a board would read it. The agreement funds USD 313,000 over a twelve-month term from 1 August 2023, payable either as twelve monthly instalments of USD 26,083.33 or four quarterly tranches of USD 78,250, and it resolves into exactly three pillars: AI technology transfer and training, delivered as six live sessions (three online, three onsite) archived on a learning-management system and climbing from AI essentials to hands-on work with the client's own transformer model; run-manage-operate for the on-premise Nvidia DGX A100 and the RADIX MLOps stack; and continuing AI engineering. The load-bearing structural fact is that the agreement is written on an effort obligation rather than a result obligation, with a 60-day termination notice, a 30-day cure, and a liability cap set at the phase amount. We argue that this is not incidental legal boilerplate but the mechanism of the handover: a vendor paid for effort, not for a held result, has no incentive to remain indispensable, and every one of the three pillars is therefore a transfer pillar, with the training archive, the run-book, and the engineering practice all designed to be absorbed by the client's own team. Transformation succeeds, on this reading, precisely when the vendor is contractually paid to make itself replaceable. The prior chapters of this programme, the R&D funnel and the phased operating model and the people trained, are pointers here rather than re-derivations; the subject is the contract that hands the whole thing over.

Tarry Singh

Founder & CEO

Tannistha Maiti

Senior AI Researcher

TECHNOLOGY-TRANSFER / CAPABILITY-HANDOVER / SERVICES-AGREEMENT /
EFFORT-NOT-RESULT / RUN-MANAGE-OPERATE / MLOPS

CONTENTS

01	The year-two problem is not a modelling problem
02	One envelope, three pillars
03	The clause that makes the vendor removable
04	Why the archive is the pillar that compounds
05	What a board should actually check
06	The counterintuitive economics
07	Limitations
08	References

Most accounts of an AI engagement end at the moment the model works. The demonstration lands, the accuracy numbers clear the bar, the operator's team applauds, and the write-up closes on the deployment. The document that actually decides the outcome is the one nobody publishes: the services agreement for the year that follows. That agreement is where an operator learns whether it bought a capability or rented a dependency, and the two look identical on the day the model ships.

This whitepaper reads one such agreement. It is the year-two services contract from a subsurface-AI programme we ran with a mid-sized Middle East carbonate operator, and it covers the twelve months after the research year that produced the models. We read it the way a board would: not as a legal artefact to be filed, but as the design of a handover. The agreement funds USD 313,000 over a twelve-month term commencing 1 August 2023. It resolves into exactly three pillars. And it is written, deliberately, so that the vendor is paid to make itself unnecessary.

That last sentence is the whole argument, so it is worth stating plainly before any of the mechanics. A transformation programme succeeds when the vendor is contractually paid to become replaceable. Not encouraged to, not asked nicely to in a statement of work, but paid, on an effort basis, with the money attached to the act of transferring capability out of the vendor and into the client team. Everything below is the anatomy of how one agreement did that.

The year-two problem is not a modelling problem

By the time this agreement was signed, the hard machine-learning work was behind the programme. The research year had run its innovation funnel, converged on a detection-transformer lineage for fracture and bedding picking, and industrialised the result into automated tools on an on-premise control plane. We have written that story elsewhere and will not re-derive it here; the near-neighbour piece on building in-house subsurface-AI capability traces the funnel and the fifty-five people trained, and the phased operating-model piece sets out the assess-pilot-scale-operate sequence the whole engagement ran inside. This paper starts where those end.

The year-two problem is different in kind. The model exists. The question is whether the client's own engineers can run it, retrain it on next year's wells, debug it at three in the morning when a scan comes in with an unfamiliar value range, and extend it when the

geoscience team asks for a feature the vendor never built. None of that is a modelling question. All of it is a question about where the judgement lives, and judgement does not transfer because a contract says the word "handover." It transfers because someone is paid to spend a year moving it, and because the payment is structured so that the vendor's interest and the client's interest point the same way.

An agreement can get this exactly backwards without anyone noticing. A support contract that pays the vendor to keep the system running, indefinitely, on the vendor's own staff, reads as responsible stewardship and is in fact the opposite: it is a subscription to a dependency, priced to renew. The tell is not in the technology. It is in who the contract is written to make removable.

The distinction has teeth in a subsurface setting specifically because of where the models sit. This is not a cloud service the client consumes through an interface someone else operates. The platform runs on the operator's own hardware, inside the operator's own perimeter, on data that never leaves it. That deployment choice is the right one for a national operator holding data it cannot expose, but it means the client cannot outsource operation to a vendor cloud even if it wanted to. Someone inside the operator has to be able to run the machine. The only question the year-two agreement settles is whether that someone is trained during the year the vendor is paid to be there, or discovered, unprepared, the first time the vendor is not.

A board that has lived through an enterprise-software rollout will recognise the shape of the failure mode. The system goes in, the vendor's engineers keep it alive, the client's team never quite gets the keys, and three years later the renewal conversation is a hostage negotiation because nobody on the client side can do the work and everybody knows it. The technology can be excellent and the outcome still be a dependency. What prevents that outcome is not better technology. It is a contract whose funded work is the transfer itself.

One envelope, three pillars

The agreement funds a single envelope, USD 313,000, and scopes three service pillars against it. The cadence is a board's first question and the least interesting answer: the same envelope flows either as twelve monthly instalments of USD 26,083.33 or as four

quarterly tranches of USD 78,250, invoiced across the twelve-month term. Both cadences are the same total divided the same way, the envelope spread evenly across the term:

MONTHLY AND QUARTERLY INSTALMENTS ARE THE SAME ENVELOPE
OVER THE SAME TERM

Formula

LaTeX

Copy LaTeX

$$\frac{\$313,000}{12 \text{ months}} = \$26,083.33 = \frac{\$313,000}{4 \text{ quarters}} = \$78,250$$

The invoice cadence is a treasury preference. It changes nothing about what the money buys.

What the money buys is three things, and the shape of the three is the point.

Pillar A, AI technology transfer and training. Six live sessions, three delivered online and three onsite at the client's premises, archived on a learning-management system. The curriculum is a climb: AI essentials and a shared vocabulary first, then the data-engineering pipeline that ingests and normalises the logs, then AI applied to the geoscience itself (vugs, fractures and beddings, seismic machine learning), then a transformer deep-dive conducted hands-on with the client's own model, then integration and deployment, and finally a two-part MLOps track on running and maintaining the system. This is transfer by construction: the deliverable is not a document, it is a set of people who can do the work.

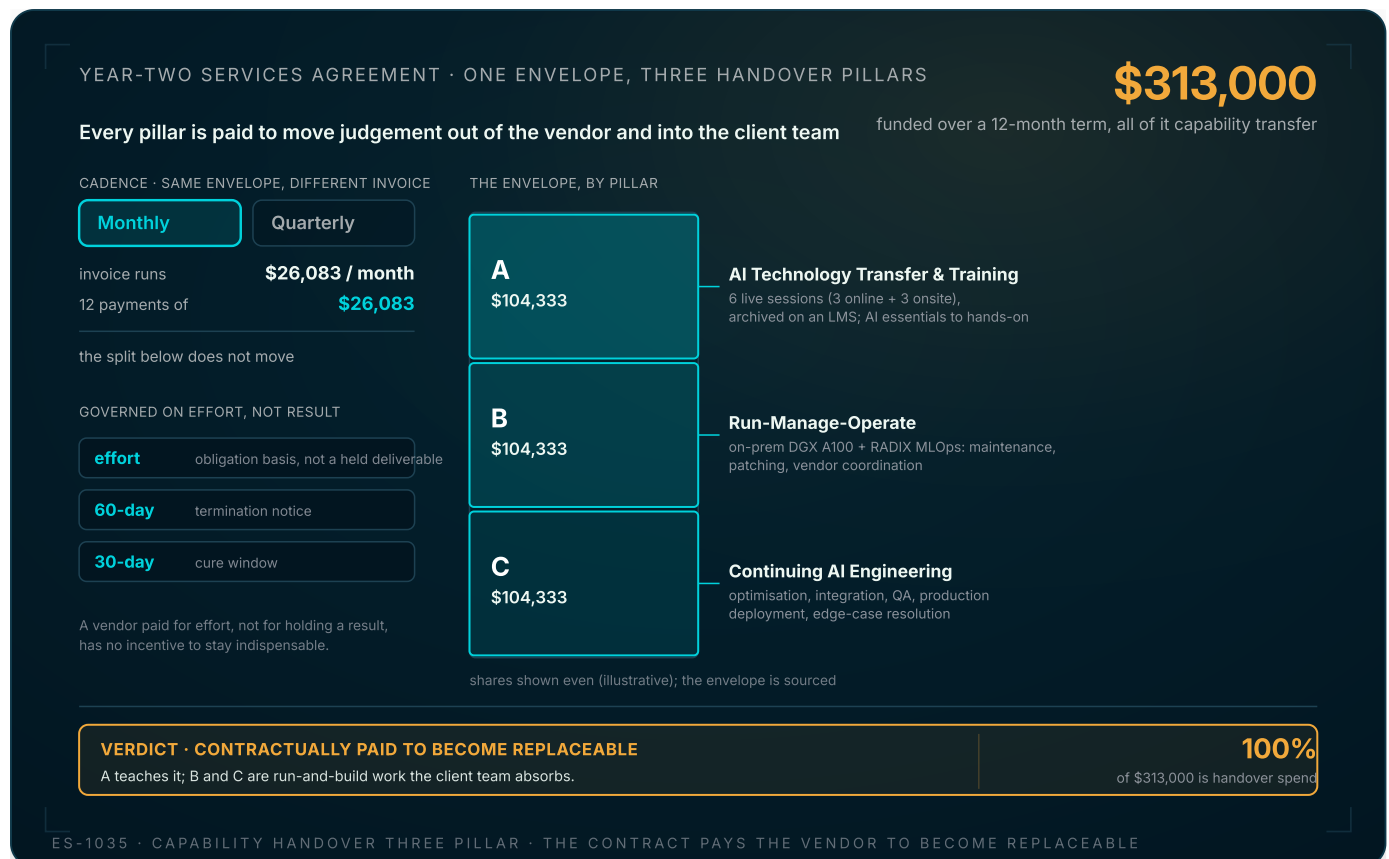
Pillar B, run-manage-operate. Maintenance of the on-premise Nvidia DGX A100 systems and the RADIX MLOps tooling, operating-system and library patching, and coordination with hardware and software vendors. This is the unglamorous middle of any AI operation, the part that has nothing to do with model accuracy and everything to do with whether the platform is up on a Tuesday. It is scoped as a service the client's own operations team is meant to absorb, with the vendor running it while that team learns to.

Boards routinely underweight this pillar, and it is worth being concrete about what it contains, because the concreteness is exactly what makes it hard to transfer. A DGX A100 is not a workstation the geoscience team can reboot and forget. It is a multi-GPU system with its own driver stack, its own thermal and power envelope, and a firmware and library matrix that has to be kept in step or the training jobs stop reproducing. The MLOps control plane on top of it, RADIX in this engagement, is the layer that versions datasets, tracks

experiments, and moves trained models into serving. Patching any of it is not a cosmetic update; a driver bump that is out of step with the framework version can silently change numerical behaviour, and a model that trained yesterday can fail to reproduce today for reasons that take an operations engineer, not a data scientist, to diagnose. Coordinating that across the hardware vendor, the MLOps-tool vendor, and the client's own IT is a standing discipline. Pillar B funds a year of the vendor doing that work in the open, run-book in hand, with the client's operations staff shadowing until the run-book is theirs. The deliverable of Pillar B is not uptime. It is a client team that can hold uptime after the vendor stops.

Pillar C, continuing AI engineering. Model optimisation, integration into the existing environment, quality assurance, production deployment, scaling, and edge-case resolution. This is the pillar that keeps the models honest as the data drifts and the requirements grow, and it is the one most easily mistaken for permanent vendor work. It is not scoped that way. It is scoped as the practice the client's engineers take over.

The edge cases are where this pillar earns its place, and they are specific to the geoscience. A subsurface-AI model built on one set of wells meets, in production, an image log acquired on a slightly different tool, or a well whose static value range sits outside the band the normalisation was tuned for, or an interval so heavily fractured that the picks crowd in a way the training set never showed. Each of these is a moment where the model's output has to be questioned by someone who understands both the machine learning and the geology, and each is a moment that recurs as new wells arrive. Pillar C funds the vendor working those cases alongside the client's engineers, in the actual system, so that the judgement of when to trust a pick and when to retrain migrates from the vendor's heads into the client's. Quality assurance here is not a checkbox at release; it is the ongoing practice of deciding whether a model that worked last quarter still works on this quarter's data, and that practice is precisely the thing a handover has to move.



The year-two services agreement read as a capability-handover instrument. One funded envelope, USD 313,000 over a 12-month term, splits into exactly three pillars: A, AI technology transfer and training, six live sessions (three online and three onsite) archived on an LMS, curriculum from AI essentials through hands-on work with the client's own transformer; B, run-manage-operate for the on-prem Nvidia DGX A100 and RADIX MLOps, maintenance, patching, and vendor coordination; C, continuing AI engineering, optimisation, integration, QA, production deployment, and edge-case resolution. The cadence toggle shows the same envelope either as twelve monthly payments of USD 26,083.33 or four quarterly tranches of USD 78,250; the invoice changes, the pillar split does not. The orange verdict rail carries the single argument: all three pillars move capability out of the vendor and into the client team, and the contract governs them on an effort basis, not a result basis, with a 60-day termination notice and a 30-day cure, so the vendor is paid to make itself replaceable rather than to hold a deliverable hostage. Sourced from the engagement archive: the USD 313,000 total, the 12-month term from 1 August 2023, the monthly and quarterly figures, the three pillars and their scopes, the six sessions archived on an LMS, and the effort-not-result and termination terms. The even three-way split across pillars is illustrative for shape only; the agreement prices the envelope and scopes the pillars, it does not publish a line item per pillar.

Read the three pillars together and a pattern falls out that a single-line-item budget would hide. There is no pillar for "keep the vendor's model running so you don't have to learn how." Every pillar is a transfer pillar. Pillar A transfers by teaching. Pillars B and C are run-and-build work that the client's own team is meant to take over, with the vendor present to hand it across rather than to hold it. The envelope is not a maintenance retainer wearing a handover label. It is a handover wearing the ordinary clothes of a services agreement.

The clause that makes the vendor removable

A pattern in the pillars is suggestive. It is not proof. A vendor could scope three transfer pillars and still write the governing terms to keep itself indispensable, and the terms would quietly win. So the load-bearing question is not what the pillars are called. It is how the agreement is governed.

The agreement is written on an **effort obligation, not a result obligation**. The text is explicit that the service provider is not responsible for achieving a specific result, framing the work as a research-and-engineering engagement rather than a guaranteed outcome. That single choice is what makes the handover credible rather than rhetorical, and it is worth being precise about why.

Under a result obligation, the vendor owns an outcome the client has not been taught to reproduce. Every incentive that follows tilts toward keeping the vendor in the building. Exit forfeits the promised result, so switching cost is high. Exclusivity becomes the natural hedge, because a vendor on the hook for a result wants to control every input to it. Liability negotiations turn adversarial, because the cap sits under an outcome nobody can fully guarantee. The contract becomes a machine for lock-in, and it does so through terms that each look reasonable in isolation.

Under an effort obligation, the same terms invert. The vendor is paid to teach, to run, and to build, and there is no held result to forfeit on exit. So the 60-day termination notice and the 30-day cure window make the client genuinely free to leave mid-transfer rather than trapped by it. The liability cap set at the phase amount fits paid-for effort rather than sitting under an unguaranteeable outcome. The non-exclusivity clause lets the client staff up and second-source freely while the transfer is under way, which is exactly what a client absorbing a capability needs to be allowed to do. Every governance term that would tilt toward lock-in under a result obligation instead tilts toward handover.

A vendor paid for effort, not for a held result, has no reason to stay indispensable

LEVER · OBLIGATION BASIS

Effort (sourced)

Result (counterfactual)

Obligation basis

sourced: effort, not a specific result

vendor paid to teach, run, build - no outcome to withhold

60-day notice / 30-day cure

sourced: either party exits on notice

exit is cheap; the client is never trapped mid-transfer

Liability capped at phase

sourced: cap = the phase amount only

no result to indemnify; the cap fits paid-for effort

Non-exclusivity

sourced: neither side is locked in

client may staff and second-source freely during transfer

EFFORT OBLIGATION (THE REAL CLAUSE)

The vendor is paid to move judgement into the client team.
Every other term then makes the vendor safe to remove.

tilts toward HANDOVER

The agreement chose EFFORT. Handover was the design goal,
so the terms were written to let the vendor leave.

Clause texts and the effort basis are sourced from the agreement; the RESULT column is a labelled counterfactual, not a term of this contract.

ES-2715 · EFFORT NOT RESULT LEDGER · THE CLAUSE THAT MAKES THE VENDOR REMOVABLE

The governance half of the argument. Four sourced clauses from the year-two agreement sit on the left; a single lever flips their combined incentive between the real effort obligation and a labelled result-obligation counterfactual. Under the effort basis the vendor is paid to teach, run, and build, with a 60-day termination notice, a 30-day cure, a liability cap set at the phase amount, and a non-exclusivity clause, and every one of those terms makes the vendor cheap to remove. Under a result obligation the same terms invert: the vendor owns an outcome the client cannot reproduce, exit forfeits it, and exclusivity becomes the natural hedge, so the contract tilts toward lock-in. The orange tilt bar carries the single claim: the agreement chose effort, which is what makes the handover credible rather than rhetorical. The clause texts and the effort-not-result basis are sourced from the engagement archive; the result column is a labelled counterfactual drawn only to make the contrast legible, not a term of this contract.

The ledger above is the argument in its governance form. The effort clause is not incidental boilerplate that happened to be in the template. It is the keystone, and the surrounding terms, notice, cure, cap, non-exclusivity, all draw their direction from it. Pull the keystone and the arch inverts: the identical clauses, read under a result obligation, would describe a lock-in. The agreement chose effort, and in choosing effort it chose to be the kind of contract a vendor can be paid under and still walk away from.

There is a version of this that sounds like the vendor giving something up, and it is worth naming that it is not. An effort obligation is not weaker for the vendor; it is honest about what an AI engagement can promise. A model's field accuracy depends on data the vendor does not control and cannot guarantee, so a result obligation on model performance would

be a promise written on someone else's inputs. The effort basis is the accurate description of the work, and it happens to be the basis that also makes the handover possible. The interests line up because the contract is telling the truth about the engagement.

The intellectual-property terms carry the same logic and deserve their own line, because IP is where lock-in most often hides in AI contracts. This agreement keeps each party's intellectual property with that party and forbids use without written consent, with foreground IP defined so that what the engagement produces has a named owner rather than an ambiguous claim the vendor can later assert. That matters for the handover in a way that is easy to miss. A contract that quietly vests the trained models, the pipelines, or the tooling in the vendor has built a dependency no amount of training can dissolve, because the client can be taught to run a system it is not permitted to own. Clear IP allocation is the precondition that lets the training and the run-book and the engineering practice actually land somewhere the client controls. A handover into a system the client cannot legally keep is not a handover.


None of these clauses is unusual on its own. Notice periods, cure windows, liability caps, non-exclusivity, and IP allocation appear in almost every services agreement, which is exactly why they are easy to sign without reading them as a system. The point is not that any single term is remarkable. It is that read together, under the effort obligation, they compose into a contract that is cheap to leave and clear about ownership, and a contract that is cheap to leave and clear about ownership is the only kind under which a real handover can occur. The same terms under a result obligation compose into the opposite. The clauses are the same words; the basis clause decides what they mean.

Why the archive is the pillar that compounds

Of the three pillars, the training pillar is the one most likely to be underestimated by a board reading the budget, because training reads as a soft line item, a nice-to-have workshop series. On this reading it is the opposite: it is the pillar with the longest half-life, and the reason is a single design decision inside it.

The six sessions are archived on a learning-management system. That is the whole difference between a training event and a transferred capability. A live session, run once, transfers its content to whoever is in the room those days, and then it is gone. The people who were there carry it; the people the client hires next year do not. An archived session

transfers indefinitely. The same six sessions become a standing curriculum that every engineer the client onboards after the vendor has left can climb, from AI essentials to hands-on work with the client's own transformer, without the vendor in the room.



Pillar A up close. The technology-transfer and training pillar is six live sessions, three online and three onsite, climbing from AI essentials through data engineering, AI applied to geology, a transformer deep-dive with hands-on work on the client's own model, integration and deployment, and a two-part MLOps track. The toggle carries the argument: run the sessions live only and the retained-capability band collapses to the handful of days the sessions actually ran, a one-time transfer to whoever was in the room; archive them on an LMS and the band fills the whole climb, because the same curriculum keeps transferring capability to every engineer the client hires after the vendor has left. The orange band is the retained asset, which is the reason the sessions were archived rather than simply delivered. Sourced from the engagement archive: the six sessions, the three-online and three-onsite split, the LMS archive, and the curriculum arc from AI essentials to the client's own transformer. The per-step arrangement of online versus onsite is an illustrative reading of the sourced three-and-three counts, and the retained-capability band is a qualitative argument, not a measured duration.

The instrument above makes the compounding visible. Toggle the archive off and the retained capability collapses to the handful of days the sessions actually ran, a one-time transfer with a short memory. Toggle it on and the capability persists past the engagement, because the asset is not the delivery, it is the durable record of it. This is why the

agreement specifies that the sessions are archived rather than merely delivered. A board that reads Pillar A as "six workshops" has mispriced it. The correct reading is that Pillar A builds a piece of institutional memory that keeps paying out after the money stops.

The curriculum arc matters here too, and not only its storage. The climb ends at the client's own transformer model, hands-on, not at a generic tutorial. The final rung of the training is the client's engineers working inside the actual system they will own, which is the difference between learning about transformers and learning this transformer. Combined with the archive, that arc is the mechanism by which the deepest, most engagement-specific knowledge, the part hardest to hand over, becomes the part a new hire can pick up two years later from a recording.

What a board should actually check

The temptation, reading all this, is to treat "make yourself replaceable" as a slogan and move on. It is not a slogan; it is a checklist, and the items are concrete enough that a board can hold an engagement to them before signing the year-two agreement rather than discovering the answer a year too late.

- Does the agreement fund transfer as its own scoped work, or is transfer an unpriced promise stapled to a maintenance retainer? If there is no line of effort whose deliverable is the client's own people doing the work, there is no handover, whatever the cover page says.
- Is the governing obligation effort or result? A result obligation on model performance is a promise written on data the vendor does not control, and it quietly converts the whole relationship into lock-in. The effort basis is both the honest description and the one that keeps the client free.
- Are the exit terms cheap? A short termination notice, a real cure window, and non-exclusivity are not vendor concessions; they are the client's insurance that the transfer is happening because it works, not because they are trapped.
- Is the training archived, or merely delivered? An unarchived training pillar transfers capability once. An archived one transfers it to everyone the client hires next. The storage decision is the compounding decision.
- Does the curriculum end inside the client's own system? Generic training builds generic knowledge. The transfer that matters ends with the client's engineers hands-on in the actual platform they will own.

Every one of these is a structural fact about the contract, checkable at signing, independent of how good the model is. That independence is the point. The model's quality is the year-one question, and it is answered before this agreement is drafted. The year-two question is entirely a question of structure, and a strong model handed over on a lock-in contract is still a dependency.

The counterintuitive economics

There is a final objection worth meeting head on, because it is the one a sharp board member will raise: does a vendor not lose by writing itself out of the account? Why would any service provider design an agreement whose success condition is its own redundancy?

The answer is that the vendor is not writing itself out of the account. It is writing itself out of *this* account, on *these* three pillars, and it is doing so in a way that produces the one asset that generates the next engagement: a reference client that owns a working capability and attributes it to the transfer. A dependency does not refer; it resents, quietly, the sense of being locked in, and it churns the moment a credible alternative appears. A client that was genuinely handed a capability, on a contract that made leaving cheap and chose not to, is the most durable commercial asset a services business can hold. The effort obligation that makes the vendor replaceable on the year-two pillars is the same clause that makes the vendor the obvious partner for whatever the client wants to build next, precisely because the client was never trapped into the last thing.

So the economics are not a sacrifice. They are the recognition that in a transformation engagement, the vendor's interest and the client's interest are aligned only when the contract is written to align them, and that the alignment runs through replaceability. A vendor paid to hold a result optimises for holding it. A vendor paid for the effort of transfer optimises for the transfer, and the transfer, done honestly and archived and governed to be cheap to leave, is what a board is actually buying in year two.

There is also a sovereignty dimension that a national operator, in particular, cannot treat as a nicety. Subsurface data is strategic, and the capability to interpret it is a form of national industrial capacity, not just a vendor service line. An operator that owns its models, its run-book, and its trained people owns its own subsurface interpretation; an operator that rents them has outsourced a strategic function to a foreign vendor it cannot fully audit and cannot cheaply replace. The year-two agreement, read at that altitude, is not a compute-

services purchase. It is the instrument through which a strategic capability comes home. The effort obligation, the archived curriculum ending at the operator's own transformer, the run-book handed across on Pillar B, and the clean IP allocation are, taken together, the mechanism by which the operator ends the year able to do the work without asking anyone's permission. That is worth more than the USD 313,000 line suggests, and it is the real return a board should price when it evaluates the agreement.

The year-one engagement builds the model. The year-two agreement decides whether the operator owns it. The agreement we have read here decided in favour of the operator, and it did so not through goodwill but through structure: one funded envelope, three transfer pillars, and an effort obligation that made every governing term lean toward the door. That is what it looks like when a transformation is designed to succeed after the vendor leaves.

Limitations

This is a board-level synthesis of one year-two services agreement from a single subsurface-AI programme, anonymised, and it should be read as an argument about contract structure rather than as a legal template or a generalisable benchmark. Three cautions in particular.

First, the per-pillar split of the USD 313,000 envelope is illustrative. The agreement prices the envelope and scopes the three pillars; it does not publish a line item per pillar, so the even three-way division used in the primary instrument is a reading for shape only, not a sourced allocation, and the real internal weighting of the three pillars could differ substantially.

Second, the result-obligation column in the governance ledger is a labelled counterfactual, drawn to make the effort clause legible by contrast. It is not a term of this agreement and not a claim about any specific competing contract; it is the analytic other-pole against which the sourced effort basis is read.

Third, the claim that an effort-and-transfer structure produces better long-run outcomes than a result-and-retainer structure is an argument grounded in this engagement's design and our experience across engagements, not a controlled comparison. We have not run, and could not ethically run, the same programme twice under opposing contract structures

to measure the difference. What the evidence rests on is the internal consistency of one agreement whose stated purpose, funded pillars, and governing clauses all point the same way, toward a vendor paid to become replaceable.

References

Koroteev, D., Tekic, Z. (2021). Artificial intelligence in oil and gas upstream: Trends, challenges, and scenarios for the future. *Energy and AI*, 3, 100041. The upstream survey framing why an operator carries subsurface-AI capability in-house rather than renting it per project. <https://doi.org/10.1016/j.egyai.2020.100041>

Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., et al. (2017). Artificial Intelligence: The Next Digital Frontier? McKinsey Global Institute. On the organisational, not merely technical, conditions under which enterprise AI adoption durably sticks. <https://www.mckinsey.com/~/media/mckinsey/industries/advanced%20electronics/our%20insights/how%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/mgi-artificial-intelligence-discussion-paper.pdf>

After the Model Ships: Anatomy of a Year-Two AI Technology-Transfer Engagement

Published December 2025. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/anatomy-of-a-year-two-ai-technology-transfer-engagement>

Copyright EarthScan. All rights reserved.