

The Contract Design Behind a Successful AI Capability Transfer

Most accounts of an AI engagement are told through the model. The document that decides whether the operator ends up owning a capability or renting a dependency is the contract, and on a programme that runs across years it is not one contract but two: the founding research-service agreement and the year-two services agreement that follows. We read both, from a subsurface-AI programme built for a major operator in Oman on borehole-image logs from a fractured carbonate reservoir, as one design discipline whose object is a clean exit. Four levers carry that discipline. The obligation is written as effort, not a warranted result, on genuinely uncertain research. The foreground intellectual property is split 51 to the operator and 49 to the vendor, the smallest majority that still hands the operator control while leaving the vendor free to keep improving its own methods. A funded year-two agreement, USD 313,000 over twelve months, buys a handover rather than a subscription: six live training sessions, three online and three onsite, archived on a learning system, plus run-manage-operate and continuing engineering. And the money moves on a cadence that rewards transfer, five phase-gated tranches summing to USD 345,581.60 in year one, then a flat run-out in year two. A fifth structural move keeps the design clean: the academic partner's research is carved to a separate agreement so it does not entangle the operator's holdings. The two prior chapters of this programme, one reading the year-one paperwork as risk transfer and one reading the year-two agreement as a handover, are pointers here rather than re-derivations; the subject of this piece is the doctrine that spans both, and the argument that a vendor which designs its own replaceability into the contract is the one worth signing. Three interactive instruments make the case concrete.

Tarry Singh

Founder & CEO

Tannistha Maiti

Senior AI Researcher

CONTRACT-DESIGN / CAPABILITY-TRANSFER / VENDOR-REPLACEABILITY / EFFORT-OBLIGATION / FOREGROUND-IP / SELF-SUFFICIENCY

CONTENTS

01	The doctrine on one board
02	Effort is the clause that makes replaceability affordable
03	The IP split hands over the majority, on purpose
04	The cadence pays for transfer, not for tenure
05	What the operator holds when the vendor leaves
06	The fifth move: keeping the transfer terms clean
07	Why a vendor would design its own exit
08	What a board should read before it signs
09	Limitations
10	References

A services vendor with a strong hand usually writes the contract to keep its hand strong. It warrants a result the buyer will pay to protect, it holds the intellectual property so the buyer cannot leave, and it prices the relationship so that next year's invoice is the natural continuation of this year's. Every one of those moves is legible, defensible, and, on an AI capability programme, wrong. They produce a client that is locked in and quietly resentful, and a capability that lives in the vendor rather than in the operator who paid for it.

This whitepaper is about the opposite design. We ran the technical side of a subsurface-AI programme for a major operator in Oman, on borehole-image logs from a fractured carbonate reservoir, and it was governed not by one contract but by two: a founding research-service agreement for the year of model-building, and a year-two services agreement for the year of handover. Read together, the two documents are a single discipline whose object is a clean exit. The vendor is engaged, deliberately, to make itself replaceable. Not encouraged to in a statement of work, but structured to, clause by clause, with the money attached to the act of moving capability out of the vendor and into the operator's own team.


We have written the two halves of this story before, and this piece stands on both rather than repeating either. One earlier whitepaper reads the year-one agreement as a risk-transfer machine, five levers moving research risk, liability, funding timing, exit rights, and IP between the parties. Another reads the year-two agreement as a handover, USD 313,000 resolved into three transfer pillars. This paper sits above both. Its subject is the design doctrine that spans the two years, and its claim is a counter-intuitive one that a board should test rather than assume: on an AI capability programme, the vendor most worth signing is the one that has written its own obsolescence into the paperwork.

The doctrine on one board

Strip the two agreements down to the clauses that decide who owns the capability at the end, and a small number of levers do all the work. Each is an ordinary contract term. What makes them a doctrine rather than a checklist is that they all lean the same way, toward the operator being able to run, own, and staff the capability without the vendor.

There are four, and they recur across both years. The obligation is written as effort, not a warranted result. The foreground intellectual property is split so the operator holds the majority. A funded year-two agreement buys a handover rather than a subscription. And the

payment cadence pays the vendor for building the operator's self-sufficiency rather than for holding a dependency. A fifth move sits underneath them and keeps the design clean: the academic partner's research is carved into a separate agreement, so the transfer terms between operator and vendor do not have to carry the university's own interests.



The two signed agreements behind one subsurface-AI programme, read as a single design discipline: every governing clause is shaped to push capability out of the vendor and into the operator. Four exit-by-design levers sit as rows, each carrying its sourced term and a teal arrow whose length shows how hard that lever leans the operator toward doing the work alone; every arrow points at the one orange element, the self-sufficiency verdict rail they converge on. Toggle between the Year-1 research-service draft (15 Nov 2021, USD 345,581.60, 2-year term, 51/49 foreground-IP split to the operator) and the Year-2 technology-transfer agreement (1 Jul 2023, USD 313,000, 12-month term, six live training sessions archived on a learning system); the sourced terms on two levers change, the design intent does not. A fifth structural move sits below the board: the academic partner's research is carved to a separate agreement so the transfer terms stay clean. The clause terms, totals, term lengths, IP split, exit notice, and training-session counts are sourced from the engagement archive; the 0-to-100 self-sufficiency pull scores are an illustrative reading of design intent, not a measurement.

The board above lays the four levers side by side and lets you toggle between the two agreements to see which terms change and which hold. The sourced numbers move, the total drops from USD 345,581.60 in the research year to USD 313,000 in the transfer year, the two-year term contracts to twelve months, the IP clause shifts from a 51-49 split to a

each-keeps-their-own arrangement once the foreground already sits with the operator. The design intent does not move at all. In both years every governing lever points at the same right-hand rail: the operator ends the engagement able to do the work alone. The rest of this piece is an account of why each lever is shaped that way, and what it would cost to shape it the other.

Effort is the clause that makes replaceability affordable

The load-bearing lever is the one that reads like a disclaimer. In both agreements the vendor is engaged on an effort basis and is, in the contract's own words, not responsible for achieving a specific result, because the work is a research project. The year-one whitepaper works the arithmetic of that clause in full, why a results warranty on genuinely uncertain research is an insurance policy the vendor cannot price, and why the effort framing is what lets the research be bought at all. We will not re-derive that here.

What matters for the doctrine is a second, quieter job the same clause does, and it is the job that makes replaceability affordable. A vendor paid to hold a warranted result has every reason to stay indispensable, because its exposure only ends when it can prove the result holds, and the surest way to keep proving it is to keep running the system itself. A vendor paid for effort has no such incentive. It is paid for the work of transfer, and the work of transfer is finished when the operator's team can carry the result forward without it. The effort clause does not merely make the research signable in year one. In year two it removes the vendor's reason to cling. A vendor that is never on the hook for a held result is a vendor that can be paid, in good conscience, to work itself out of the account.

That is why the same clause appears in both agreements doing two different jobs. In the research year it lets the operator buy an uncertain thing without asking the vendor to insure the outcome. In the transfer year it lets the vendor hand the thing over without breaching a warranty it would otherwise have to keep alive. One clause, read across two years, is the hinge the whole doctrine turns on.

It is worth being precise about what the effort clause does not do, because the objection to it is usually that it lets a vendor coast. It does not. Both agreements carry a defined scope, a phase structure with deliverables, a governance cadence, and the ordinary remedies for a vendor that simply fails to do the work. The effort framing draws a line between two very different failures: the failure to pursue the work competently, which remains squarely the

vendor's problem, and the failure of a competently-pursued research question to yield the hoped-for answer, which is the operator's, because the operator chose the question and owns the upside if it lands. A results warranty erases that line and bills the vendor for both sides of it. The effort clause preserves it, and preserving it is what keeps the price honest and the exit clean. A vendor that had sold a warranty would have to stay attached to the result to defend it; a vendor that sold effort against a scope has discharged its obligation the moment the scope is delivered and the capability is in the operator's hands.

The IP split hands over the majority, on purpose

The second lever is ownership, and its shape is the most deliberate number in the year-one agreement. The foreground intellectual property, everything the work newly creates, is split 51% to the operator and 49% to the vendor and its company jointly. 51 is not a round number chosen for symbolism. It is the smallest majority that still hands the operator control, and the choice to stop there is the whole argument.

A vendor optimising for lock-in keeps the IP. A vendor optimising for a quick sale gives all of it away and loses the right to build on its own methods. The 51-49 split does neither. It gives the operator the controlling stake in what its money created, so the operator owns the capability rather than licensing it, while leaving the vendor a real, minority interest and, more importantly, the freedom to keep improving the underlying methods on the next engagement. The operator cannot be held hostage over IP it majority-owns. The vendor is not stripped of the techniques that make it worth hiring again. Both of those are conditions for a clean exit: the operator leaves owning what it paid for, and the vendor leaves with its craft intact.

By the year-two agreement the foreground already sits substantially with the operator, so the IP clause there does the simpler thing, each party keeps its own and neither uses the other's without written consent. The heavy lifting was done in year one, when the split was set to a majority. That is the point in the timeline where a lock-in vendor would have written the clause to keep the operator dependent, and it is exactly the point where this design chose not to.

There is a detail in how the majority is held that matters for the exit. The operator's 51% is a controlling stake in the foreground as a whole, not a carve-up of specific artefacts, so there is no argument to be had later about which model, which script, or which dataset belongs to


whom. Control is settled at the level of the foreground, cleanly, once. The vendor's 49% is not a lever it can pull to obstruct the operator; a minority cannot block a majority's use of what it owns. What the vendor keeps is the thing it actually needs to remain a viable business, the freedom to carry its own methods, its own architectural choices, its own accumulated craft, into the next engagement without asking the operator's permission. The split is therefore not a compromise that leaves both sides half-satisfied. It is a deliberate allocation that gives each party exactly the thing that matters to it and denies each party the thing it could use to trap the other. That is a rarer piece of drafting than it sounds, and it is the reason the IP clause is a lever toward the exit rather than a snag on the way to it.

The cadence pays for transfer, not for tenure

The third and fourth levers are best read together, because they are the same idea seen from two angles: the handover is funded, and the money moves on a schedule that rewards building the operator's independence.

The funded handover is the year-two agreement itself. Rather than let the relationship drift into an open-ended support retainer, the two sides wrote a bounded, twelve-month, USD 313,000 agreement whose stated purpose is to move capability across. Its centre is a training programme delivered as six live sessions, three online and three onsite at the operator's premises, archived on a learning system so the curriculum survives the people who first sat through it. Around that sit a run-manage-operate pillar for the on-premise compute and MLOps stack, and a continuing-engineering pillar. The year-two whitepaper reads those three pillars in detail as a handover; the point for this doctrine is only that the handover was paid for as its own thing, with its own envelope, rather than smuggled into a support contract that would have quietly reset every year.

The cadence is where the design shows its hand most clearly.



The full two-year payment schedule, laid on one timeline and read as a transfer plan. Year 1 releases five phase-gated tranches (84,047.60 / 66,648.40 / 64,584.00 / 78,717.60 / 51,584.00, summing to USD 345,581.60, invoiced every five months from Nov'21 to Jul'23); Year 2 pays USD 313,000 flat over a twelve-month handover term, shown here as the four quarterly tranches of USD 78,250 (Oct'23 to Jul'24), the shape of a run-out rather than a lock-in. The teal step curve is cumulative USD released; drag the orange handover scrubber and the marker reads the one number that argues, the money released at that point in the transfer, so money and capability visibly move together. The tranche amounts, dates, totals, invoice cadence, and the monthly/quarterly Year-2 split are sourced from the engagement archive; the transfer-milestone label on each tranche is an illustrative reading of what the money bought at that gate, not a sourced line item, and the timeline positions are drawn to scale the two agreements against each other, not to a precise calendar.

Laid on one timeline, the two agreements' payment schedules read as a transfer plan. The research year releases five phase-gated tranches, 84,047.60, then 66,648.40, 64,584.00, 78,717.60, and 51,584.00, summing to USD 345,581.60 and invoiced every five months from late 2021 to mid-2023. Each tranche is gated to a phase completing, and each phase moves something concrete closer to the operator: the infrastructure stood up, the dataset built with the operator's own data, the models trained, the tools industrialised, the year-one capability handed across. Then the year-two envelope pays flat, USD 313,000 over twelve months as either monthly instalments or four quarterly tranches of USD 78,250. Drag the scrubber and the shape of the deal is legible: the money released tracks the capability moved, and the year-two run is a flat run-out, not the accelerating curve of a subscription

that grows the longer the client stays. A cadence built for tenure would back-load, or renew, or price the later work higher to reward incumbency. This one does the opposite. It pays most heavily for the phases that build and transfer, and it flattens toward zero-slope in the year whose job is to leave.

We can state the year-two envelope's neutrality between its two payment options directly, because it is the clearest small sign that the cadence is a run-out and not a growth curve:

THE YEAR-TWO ENVELOPE IS THE SAME WHETHER PAID MONTHLY OR
QUARTERLY

Formula

LaTeX

Copy LaTeX


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

Monthly or quarterly, the total is fixed and the term is fixed. There is no mechanism in the year-two agreement by which staying longer costs the operator more per unit of capability, and no clause that makes the vendor richer for remaining indispensable. The cadence is designed to be spent and finished.

The year-one cadence carries a second property that reinforces the same intent, and it is one the risk-transfer whitepaper reads in full: the liability the vendor carries is capped at the value of the phase in which a claim arises, and never reaches back to the fees for phases already completed and paid. The effect on the exit is that completing a phase genuinely retires the risk attached to it. Each of the five tranches is not just a payment but the closing of a self-contained unit of risk, so by the time the research year ends there is no accumulated, contract-wide exposure hanging over either party that would make them reluctant to part. A flat cap at total contract value would have done the opposite, leaving a dispute in the last phase able to claw back the whole programme's fees and giving the vendor a reason to stay entangled in order to defend them. The per-phase staircase clears the ground for a clean handover: each phase is finished, paid, and closed on its own terms, and nothing about the earlier work remains open to reopen the relationship.

What the operator holds when the vendor leaves

The four levers are means. The end they serve is what the operator actually owns on the day the engagement closes, and that is the cleanest way to test whether the design worked.



What the operator actually holds the day the engagement ends, drawn as an endowment rather than an expense. Under the exit-by-design contracts, five assets land on the operator's side of the ledger as things it owns and can run without the vendor: the trained models (delivered on an effort basis), the majority of the foreground IP (the 51/49 split to the operator in the Year-1 draft), the run-book and operations (the Year-2 run-manage-operate pillar), a trained team (the six live training sessions, three online and three onsite, archived on a learning system), and the university research ring-fenced in a separate agreement so it does not entangle the operator's holdings. Toggle to the dependency design and the same five flip to rented access that ends the day the vendor leaves; the orange verdict counts how many of the five the operator owns outright. The owned outcomes trace to sourced contract terms in the engagement archive; the dependency column is a labelled counterfactual, the analytic other-pole against which the actual contract is read, not a term of any real agreement.

Read as an endowment rather than an expense, the exit-by-design contracts leave the operator holding five things it can run without the vendor. It owns the trained models, delivered on an effort basis against its own data. It owns the majority of the foreground IP, through the 51-49 split. It owns the run-book and the operations, transferred through the

year-two run-manage-operate pillar. It owns a trained team, built through the six archived training sessions. And the university's research is ring-fenced in its own agreement, so the operator's holdings are not entangled with an academic partner's separate interests. Toggle the instrument to the dependency design and the same five assets flip to rented access that stops the day the vendor walks. That contrast is the whole return on the doctrine. The same spend, under the same programme, either buys a capability the operator keeps or an access it loses, and the difference is entirely in how the contract was written.

The dependency column is a counterfactual, not a description of any real competing agreement. It is drawn to make the owned outcomes legible by contrast, and the honest reading is that most services contracts sit somewhere between the two poles rather than at the rented extreme. But the poles are real, and the choice between them is real. A board evaluating an AI vendor is choosing, whether or not it names the choice, how close to the owned pole its own engagement will land.

The fifth move: keeping the transfer terms clean

The four levers are the visible doctrine. Underneath them sits a structural decision that makes the whole thing tractable, and it is worth naming because it is the kind of move a hurried contract skips.

The academic partner's research was carved into a separate agreement. On this programme the vendor worked alongside a university, and a university brings its own interests, publication rights, the principal investigator's standing, the scholarly output that a research group needs and that a commercial operator has no reason to fund or own. Folding all of that into the same document that governs the operator-vendor transfer would have entangled two very different deals: one about handing an operator a working capability, and one about a university doing publishable research. Kept in one contract, every transfer term would have had to account for the university's separate claims, and the clean 51-49 split between operator and vendor would have blurred.

Separating them keeps each deal honest. The operator-vendor agreement can be written purely around the transfer, with an IP split that hands the operator the majority and an effort obligation that lets the vendor leave. The university research agreement can carry the scholarly interests without dragging them into the operator's holdings. The two are related, they draw on the same underlying work, but they are ring-fenced from each other by

design. That ring-fencing is why the exit is clean: when the engagement ends, the operator's endowment is not encumbered by a third party's rights that were never meant to sit inside it.

Why a vendor would design its own exit

The obvious objection is the one a sharp board member raises immediately, and it deserves a direct answer rather than a reassuring one: why would any competent vendor design a contract whose success condition is its own redundancy? A vendor that writes itself out of the account is, on the face of it, writing itself out of revenue.

The answer is that the vendor is not writing itself out of the account. It is writing itself out of this account, on these specific deliverables, and it is doing so in the one way that produces the asset that generates the next one: a reference operator that owns a working capability and attributes it to the transfer. A locked-in client does not refer; it endures the lock-in and churns the moment a credible alternative appears. An operator that was genuinely handed a capability, on a contract that made leaving cheap and that the operator chose to renew anyway, is the most durable commercial relationship a services business can hold. The effort obligation that makes the vendor replaceable is the same clause that makes it the obvious partner for whatever the operator builds next, precisely because the operator was never trapped into the last thing.

There is a second reason, and for a national operator it outranks the first. Subsurface data is strategic, and the capability to interpret it is a form of industrial capacity, not a vendor convenience. This is where the Omanization dimension of the programme stops being a compliance line and becomes the point. The year-two training sessions, three of them onsite and all of them archived, were building local people who could run the models after the vendor left. An operator that owns its models, its run-book, and its trained team owns its own subsurface interpretation; an operator that rents them has outsourced a strategic national function to a foreign vendor it cannot fully audit and cannot cheaply replace. The exit-by-design doctrine, read at that altitude, is not a commercial nicety. It is the mechanism by which a strategic capability comes home, and it is worth more to the operator than any single line on either invoice suggests.

So the economics are not a sacrifice. They are the recognition that on a capability programme the vendor's interest and the operator's interest align 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, on a cadence that funds transfer and an IP split that hands over the majority, optimises for the transfer, and the transfer, done honestly and archived and ring-fenced to be clean to leave, is what a board is actually buying.

What a board should read before it signs

For a board or a procurement lead evaluating an AI capability vendor, the doctrine reduces to a short list of things to check in the paperwork itself, before the model demonstration flatters everyone into signing.

- Is the obligation written as effort against a defined scope, or as a warranted result? A warranty on uncertain research is either unpriceable or a hidden premium, and either way it points the vendor toward staying indispensable.
- Does the foreground IP hand the operator the majority? Anything less leaves the operator licensing what it paid to create, and anything that strips the vendor entirely tends to price the deal out.
- Is the handover funded as its own bounded agreement, or is it left to drift into a support retainer that resets every year? A funded, term-limited handover is a run-out; an open retainer is a subscription wearing a handover's name.
- Does the payment cadence reward transfer or tenure? Front-loaded, milestone-gated tranches that fund building and moving capability are the shape of an exit; back-loaded or renewing cadences are the shape of a lock-in.
- Are third-party interests, an academic partner's research, a subcontractor's methods, ring-fenced into separate agreements so the operator's endowment is not encumbered by rights that were never meant to sit inside it?

A contract that answers those five the right way is a contract designed to end well. The programme we ran answered them the right way, and the evidence is not in the model metrics but in the state the operator was left in: owning the majority of the IP, holding the run-book, staffed with people who had sat through the archived sessions, and free to leave at any time, which is exactly why it did not need to.

Limitations

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

First, the self-sufficiency pull scores in the contract-design board are an illustrative reading of design intent, not a measurement. The clause terms they sit beside, the effort obligation, the 51-49 split, the term lengths, the exit notice, and the training-session counts, are sourced from the engagement archive; the zero-to-one-hundred pulls are our reading of how hard each lever leans toward the operator, and a different reader could weight them differently without changing the direction of any of them.

Second, the transfer-milestone label attached to each payment tranche in the cadence instrument is an illustrative account of what the money bought at that gate, not a sourced line item. The tranche amounts, dates, totals, invoice cadence, and the monthly-versus-quarterly year-two split are sourced; the story of what each gate transferred is a reading, and the timeline positions are drawn to scale the two agreements against each other rather than to a precise calendar.

Third, the dependency column in the endowment instrument is a labelled counterfactual, the analytic other-pole against which the actual contract is read. It is not a term of any real competing agreement, and the claim that an exit-by-design structure produces better long-run outcomes than a lock-in 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 run, the same programme twice under opposing contract designs to measure the difference. What the evidence rests on is the internal consistency of two agreements whose obligations, ownership terms, funded handover, and payment cadence 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>

Susskind, R., Susskind, D. (2015). The Future of the Professions: How Technology Will Transform the Work of Human Experts. Oxford University Press. On why the durable value of a professional-services relationship shifts from holding expertise to transferring it.

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>

Making Yourself Replaceable: The Contract Design Behind a Successful AI Capability Transfer

Published September 2026. Generated 2026-09-08.

<https://www.earthscan.io/whitepapers/making-yourself-replaceable-contract-design-ai-capability-transfer>

Copyright EarthScan. All rights reserved.