

From 3-4 Weeks to 2 Hours per Well: The Economic Case for Automating Borehole-Image Interpretation

Borehole-image interpretation is the most expert-bound, least scalable step in formation evaluation. A senior geoscientist spends days to weeks picking bedding planes, fractures and vugs off a single high-resolution borehole image log by hand, and the queue of unanalysed wells only grows. This whitepaper makes the economic case for automating that step with a productized deep-learning pipeline, drawn from a roughly twenty-month, three-phase engagement with a mid-sized Middle East carbonate operator we partnered with — a programme scoped against 80-plus processed image logs. We separate three returns that are usually conflated: the cycle-time compression of a single well from a multi-week manual pass to a roughly two-hour reviewed, agentic pass; the productivity and accuracy dividend the operator's interpreters captured once the system was productized — well-to-well correlation lifting interpretation productivity by about 60% and accuracy by about 75%, with the fracture and vug detectors running roughly five times faster than manual picking; and the strategic return of closing the gap between the handful of wells a manual team can interpret and the full multi-well backlog the asset actually holds. We then show why these returns only persist if the system is engineered to be owned by the operator, and we put a frame around what the reclaimed capacity is worth when it is redeployed rather than simply banked. The argument is deliberately commercial: the model is a means; the interpretation throughput it unlocks is the asset.

Tarry Singh

Founder & CEO

Tannistha Maiti

Senior AI Researcher

BOREHOLE-IMAGE-INTERPRETATION / FORMATION-EVALUATION / ROI / CYCLE-TIME
/ BOREHOLE-IMAGE-LOGS / SUBSURFACE-AI

CONTENTS

01	The cost that hides in plain sight
02	Return one: cycle-time compression per well
03	Return two: the productivity and accuracy dividend
04	Return three: closing the interpretation-capacity gap
05	Why the returns only persist if the operator owns the system
06	Putting a frame around the number
07	References

A borehole image log is one of the richest pictures a well will ever produce — bedding planes, fractures, vugs and breakouts rendered at sub-millimetre detail down the wall of the hole. It is also one of the most expensive pictures to read. The reading is done by a senior geoscientist, by hand, sinusoid by sinusoid, and it does not get cheaper with scale. Every new well adds another multi-week interpretation pass to a queue that is already longer than the team can clear. This is not a tooling inconvenience. It is the rate-limiting step on how fast an operator can convert acquired data into a structural picture it can drill against — and it is the step we set out to automate.

This whitepaper is the economic case for doing so. It is written for the people who sign off the spend and own the asset economics — the CEO, the VP of Subsurface, the asset manager, the E&P economics lead — and it is grounded in a real programme: a roughly twenty-month, three-phase engagement with a mid-sized Middle East carbonate operator we partnered with, scoped against a study set of more than 80 processed and interpreted high-resolution borehole image logs. The disciplines underneath it were deep-learning model development, computer-vision algorithm work, data and training engineering, and a production MLOps layer. But the case we are making here is not technical. It is about three distinct returns that automation produced, why they are usually conflated, and what an operator actually does with the capacity it gets back.

The cost that hides in plain sight

The cost of manual borehole-image interpretation is invisible on most cost centres because it is paid in expert time, not in line items. A geoscientist does not bill the asset for "three weeks reading one image log." The well gets evaluated, the report ships, and the cost shows up only as a backlog of wells that never get the same treatment because there were never enough expert-weeks to go around.

That manual pass is slow for reasons that are structural, not avoidable with effort. Each image log covers hundreds of metres of borehole wall. The interpreter fits a sinusoid to every dipping feature, classifies it as a bed or a fracture, reads off its dip and azimuth, and does this for thousands of features per well — while compensating in their head for tool rotation and well deviation. Classical assists exist, but they are slow in their own right: a path-opening sinusoid-fitting routine of the kind used as a pre-deep-learning baseline took

on the order of two minutes to sweep a single two-metre interval, and a literature baseline for vug analysis ran at roughly five minutes per metre. At those rates, a single well is a multi-week commitment for a senior pair of eyes — and the asset has dozens of wells.

The headline of this whitepaper is the compression of exactly that pass. A manual interpretation that occupies a senior geoscientist for the better part of three to four weeks per well collapses, once the pipeline is productized, to a reviewed pass measured in hours — on the order of two hours per well, where the machine does the picking and the expert does the judgement. That is not a marginal efficiency. It is a change of more than two orders of magnitude in cycle time, and it is the number that reframes everything downstream.



The engineering layer didn't replace the geoscientist — it moved expert time off rote evaluation onto judgment. On Operator B's 47-well portfolio, per-well Well-Log-Detection evaluation fell from 14 hours to under 90 minutes (~89%), reclaiming ~600 expert hours in 8 weeks — while geoscientist review was explicitly retained for anomaly cases and final sign-off. Sweep across the portfolio: each 14-hour bar collapses to a <90-minute sliver, the reclaimed-hours total fills toward ~600, and an orange 'review retained' band rides every converted well and never shrinks to zero. All headline numbers are the whitepaper's own; the review-band height is schematic (review is retained but its hours aren't quantified).

The instrument above makes the central move visible: automation did not delete the geoscientist, it relocated them. The hours that used to go into mechanical picking — fitting sinusoids, reading dips, transcribing into a log — are reclaimed, and the expert is retained for the part of the work that actually needs them: the anomalies, the ambiguous zones, the final sign-off. A pipeline that removed the geoscientist entirely would be a worse system and a worse economic proposition. The win is that the same expert now governs many more wells than they could ever have picked by hand.

Return one: cycle-time compression per well

The first return is the simplest to state and the hardest to overstate. Reducing the interpretation of one well from weeks to hours is a greater-than-100x compression of the binding constraint on formation-evaluation throughput.

To see why the compression is so large, it helps to separate the two things the machine does faster. The first is detection — finding and parametrising every bedding plane and fracture. In our work with the Middle East operator, the fracture and bedding detectors were measured running at roughly five times the speed of manual picking once integrated, and the vug-analysis pipeline processed image logs at under thirty seconds per metre against a classical baseline of four to five minutes per metre, an eightfold speed-up at that stage alone. The second is the elimination of the transcription-and-bookkeeping tax: the expert no longer hand-enters thousands of features, applies tool-and-well-angle corrections one feature at a time, or re-checks their own arithmetic. The compounding of "the detection is faster" and "the human bookkeeping is gone" is what turns a fivefold or eightfold component speed-up into a per-well cycle time that drops from weeks to hours.

The engineering that makes this honest matters to the economics. The detector is not a generic object-detector pointed at a log; it is a set-prediction model — a detection-transformer architecture adapted for the geometry of borehole sinusoids, trained from scratch on a light backbone rather than fine-tuned from a natural-image network, because borehole images are nothing like the photographs such networks are pretrained on. The reason that detail belongs in an economic argument is that it is what keeps the two-hour pass a *reviewed* pass rather than a *redone* pass. If the machine's picks were not trustworthy, the geoscientist would re-interpret behind it and the cycle time would collapse back toward the manual baseline. The accuracy of the model is, in commercial terms, what protects the cycle-time gain.

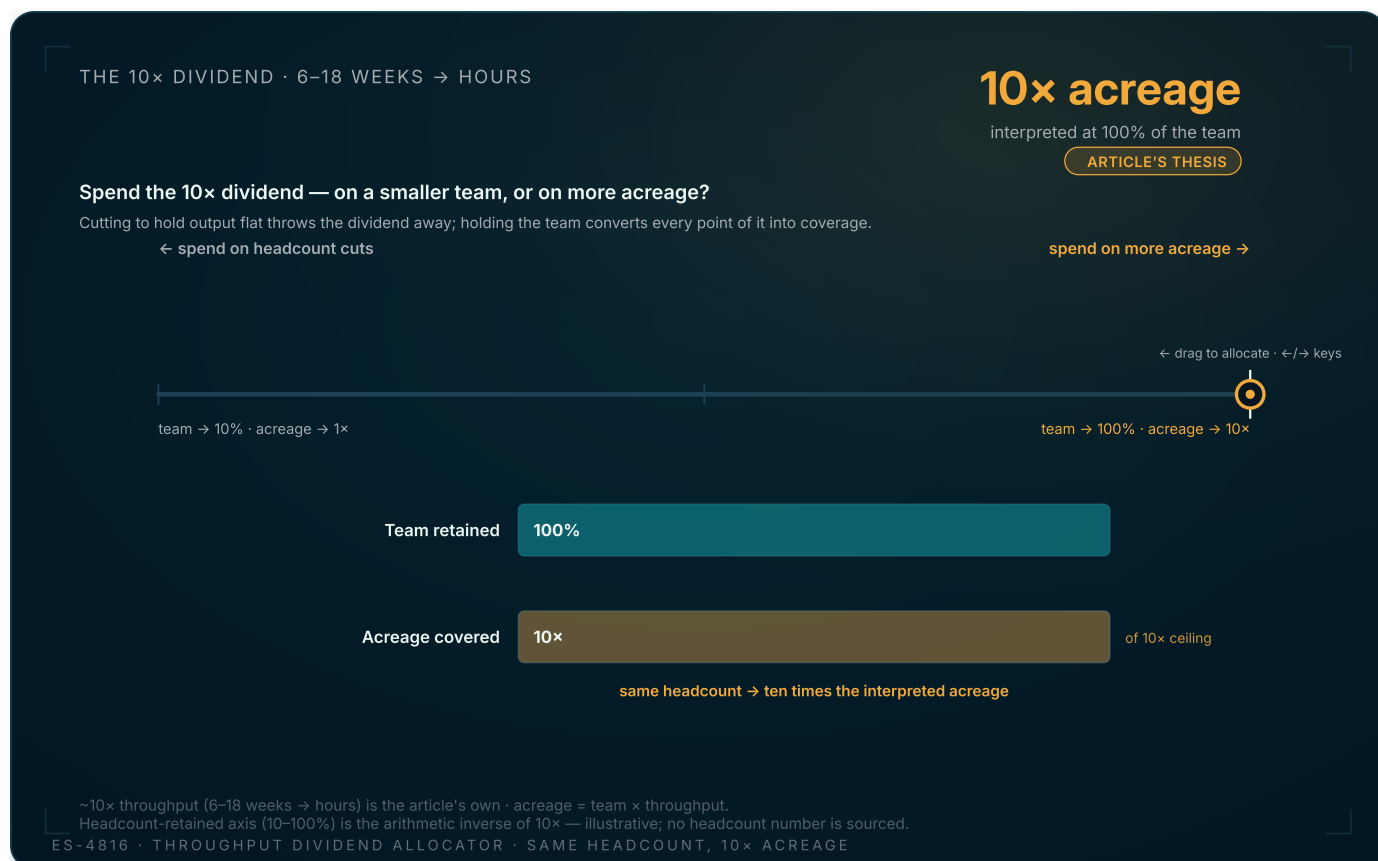
There is a unit-economics version of this that asset managers will recognise immediately. The marginal cost of interpreting the next well, under the manual regime, is roughly constant — another three-to-four expert-weeks, every time. Under the automated regime, the marginal cost of the next well falls toward the cost of a few hours of review plus a few minutes of compute. The fixed cost moved: the operator paid once, over the programme, to build and own the pipeline, and in exchange the per-well variable cost of interpretation dropped by two orders of magnitude. That is the shape of every good automation investment, and it is exactly the shape this one took.

Return two: the productivity and accuracy dividend

The second return is the one that shows up after the pipeline is no longer a research artefact and has become a tool the operator's own interpreters use day to day. It is a different number from the cycle-time headline, and conflating the two is the most common mistake made in these business cases.

Cycle-time compression is about one well in isolation. The productivity-and-accuracy dividend is about the interpreter's whole workflow once the tooling is productized. In the Middle East engagement, the well-to-well correlation capability — the part of the system that propagates a consistent structural interpretation across many wells rather than re-deriving it well by well — was measured to lift interpretation productivity by about 60% and interpretation accuracy by about 75% on a well-to-well basis, with internal precision and stratigraphic-success targets set at roughly 95% and 90% respectively. The fracture and vug detectors contributed the roughly fivefold interpretation speed-up underneath that.

Read those two numbers together and the mechanism is clear. The +60% productivity figure says the interpreter gets through substantially more work per unit time. The +75% accuracy figure says the work is also more *consistent* — the same feature interpreted the same way across wells, which is precisely what manual interpretation struggles with, because a tired human on well twelve does not pick exactly the way a fresh human picked on well one. Machine consistency is not just a quality story; it is an economic one, because inconsistent correlation between wells is what forces expensive re-work and undermines the structural model the whole asset is drilled against.




The naive read of a 10× productivity jump is 'cut 90% of the team.' The article argues the opposite: at the same headcount, the dividend buys ten times the interpreted acreage. Drag the allocator — spend the dividend on headcount cuts and acreage stays at 1×; hold the team and acreage climbs toward 10×. The single orange marker is the article's own position. The ~10× throughput multiple (6–18 weeks → hours) is sourced; the headcount-retained axis is the arithmetic inverse of 10× and is labelled illustrative, since the article states no specific headcount figure.

The allocator above is where the commercial decision actually lives, and it is the question most steering committees fail to ask. A +60% productivity uplift is not, by itself, value. It is *capacity* — and capacity is only worth something when it is deliberately redeployed. An operator can take the dividend three ways. It can bank it as cost reduction, running the same interpretation workload with fewer expert-hours. It can spend it on coverage, sending the same team across more acreage and more wells. Or it can spend it on depth, using the freed hours to interpret more thoroughly — more vug analysis, more fracture characterisation, more correlation — on the wells that matter most. The instrument lets you slide between those postures and see the trade. The point is that the productivity number is an input to a capital-allocation decision, not an output to be celebrated. The operators who get the most from automation are the ones who decide, up front, which of the three dividends they are buying.

Return three: closing the interpretation-capacity gap

The third return is strategic, and it is the one that justifies the programme rather than merely paying it back. It is the gap between how many wells a manual team can interpret and how many wells the asset actually holds.

The scoping arithmetic is stark. The study set defined for this programme ran to more than 80 processed and interpreted high-resolution borehole image logs, with the well-to-well correlation scope written to span anywhere from a handful of wells up to the full 80, and a per-phase floor of ten to fifteen wells. Against that backlog, a manual team interpreting at three-to-four weeks per well is not a team that is slightly behind. It is a team that is structurally incapable of ever catching up, because the rate at which wells are acquired exceeds the rate at which they can be read. Every drilling campaign widens the gap.



Drag the dashed analyst ceiling: the orange breach point and the “automation imperative” wedge recompute. It only flips to “within human reach” if you raise the ceiling above today’s 4,000-fold — a budget no team has. Anchors are sourced; the ceiling is an editorial throughput proxy.

The instrument above plots that divergence: interpretation demand rising with the well count, manual capacity flat against it, and the widening wedge of wells that simply never get the full-quality interpretation they warrant. This is the gap that automation closes — not

by making the team work harder, but by changing the slope of the capacity line. A pipeline that interprets a well in hours rather than weeks does not just clear the current backlog; it lifts the team's sustainable throughput above the rate of acquisition, so the wedge stops growing and starts shrinking. That is a different category of return from the first two. Cycle-time compression and the productivity dividend make the existing work cheaper and better. Closing the capacity gap makes work *possible that was never going to happen at all* — the fortieth well, the sixtieth well, the horizontal infill that would otherwise have shipped with a thinner interpretation because there was no expert-week left to give it.

It is worth being precise about scope so the number is not oversold. The validated, delivered modelling in this engagement was built and proven on a core set of 14 vertical wells logged with two different microresistivity imaging tools, later extended with five horizontal wells to validate behaviour under inclination. The 80-plus-log figure is the *backlog the productized capability is built to serve* — the scope the unlocked capacity is aimed at, not a count of wells already interpreted. We draw that line deliberately: the proven result is the pipeline and its accuracy on the validated wells; the strategic return is the backlog that pipeline makes addressable. Conflating "wells delivered" with "wells the system could serve" is exactly the kind of overclaim that erodes trust with a technical buyer, and the case is strong enough without it.

Why the returns only persist if the operator owns the system

None of these three returns survives if the system departs with the team that built it. This is the failure mode that turns a successful pilot into a dead one within a year, and it is why the economic case has to include the cost of ownership transfer rather than treating it as a closing courtesy.

The programme was structured around an explicit product-maturity arc — a research-and-development build through its early phases, with the commercial and operational stages set up as deliberate later phases rather than assumed. The build kicked off at the end of 2021 and the core delivery ran through the middle of 2023, with the commercialisation stages scoped as their own distinct effort. That sequencing is the point: productization and handover were planned line items, not afterthoughts, because a model that only the vendor can retrain is a model that delivers its returns exactly once.

COST OF INTERPRETATION, BEFORE VS AFTER

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

$$C_{\text{manual}} = N \cdot c_{\text{week}} \cdot w \quad \longrightarrow \quad C_{\text{auto}} = C_{\text{build}} + N \cdot c_{\text{review}}$$

The arithmetic above is the whole economic argument in two terms. Under manual interpretation, cost scales linearly with the number of wells, because each well costs the same multi-week block of expert time. Under the automated regime, the operator pays a one-time build-and-handover cost and then a small per-well review cost — so as the well count grows, the gap between the two curves widens without bound. But that second curve is only real if the operator can actually run the pipeline themselves. If retraining requires the vendor in the room, the build cost was a rental, not an asset, and the per-well term quietly reverts toward the manual baseline.

This is why we costed and handed over a genuine choice of operating posture rather than assuming one. The operator could run the system fully self-operated, retraining on their own hardware on a cadence of days; run it under the existing managed arrangement, retraining in minutes to hours; or push it off-premise to a managed service retraining on a two-to-three-week cadence. Three honest scenarios, with their real trade-offs, so the operator set the dial. And critically, the handover transferred judgement, not just files: the programme trained a cohort of 55 professionals, 15 of them local nationals, drawn from local universities, so the capability to run, debug and extend the platform lived inside the region rather than leaving with the delivery team. A productivity gain that evaporates when the consultants leave was never a gain. It was a loan — and the entire ownership-transfer discipline exists to make sure the operator keeps the three returns above instead of returning them.

Putting a frame around the number

We will not invent a single headline ROI figure, because an honest one depends on the operator's own loaded cost of a senior geoscientist, their well count, and which of the three dividend postures they choose. But the structure of the return is unambiguous, and it composes.

Start with cycle time: weeks to hours per well, a greater-than-100x compression of the binding constraint. Layer on the productivity-and-accuracy dividend: roughly +60% throughput and +75% well-to-well consistency once the tooling is in the interpreter's hands, underpinned by detectors running about five times faster than manual picking. Then add the strategic return that neither of the first two captures: a backlog of 80-plus wells that goes from structurally unreachable to addressable, because the capacity line now rises faster than acquisition. The first return makes each well cheaper. The second makes the team's whole output larger and more consistent. The third makes work happen that would otherwise never have happened. They are additive, and they are paid for by a one-time, ownable build whose marginal per-well cost trends toward a few hours of review.

For the executive deciding whether to fund this, the questions are not about model architecture. They are about whether the programme is structured to deliver — and keep delivering — those three returns:

- Does the business case separate the per-well cycle-time compression from the workflow-wide productivity dividend, rather than quoting one number and implying the other?
- Is the productivity uplift treated as capacity to be *allocated* — to cost, coverage, or depth — rather than a metric to be banked?
- Is the strategic return framed honestly as the backlog the capability can serve, distinct from the wells already validated?
- Is ownership transfer — operating posture, retraining cadence, and a trained local cohort — scoped and costed as a deliverable, so the returns survive the delivery team's departure?
- Is the model's accuracy understood as the thing that protects the cycle-time gain, by keeping the fast pass a reviewed pass rather than a redone one?

If the answer to all five is yes, the operator is buying a durable change in its formation-evaluation economics. If the answer to any is no, it is buying a fast demo and a slow return to the status quo.

THE ECONOMIC CASE IN FIVE LINES

- 1 Automating high-resolution borehole image-log interpretation compresses one well from a multi-week manual pass to a ~2-hour reviewed, agentic pass — a >100x cut in the binding constraint on formation-evaluation throughput.
- 2 That cycle-time headline is distinct from the productivity dividend: once productized, well-to-well correlation lifted interpretation productivity ~+60% and accuracy ~+75%, with detectors running ~5x faster than manual picking.
- 3 A +60% productivity uplift is capacity, not value — it only pays off when deliberately allocated to cost reduction, more coverage, or deeper interpretation.
- 4 The strategic return is closing the gap between the wells a manual team can read and the 80+-log backlog the asset holds — making work possible that was never going to happen at all.
- 5 All three returns are paid for by a one-time, ownable build — but they only persist if ownership (operating posture, retraining cadence, a trained local cohort) is handed over, not rented.

References

International Energy Agency, 2025. *Energy and AI Special Report*. Missing internal expertise identified as the dominant adoption barrier across the energy sector — the constraint that ownership-transfer discipline is designed to remove. <https://www.iea.org/reports/energy-and-ai>

McKinsey & Company, 2025. *The State of AI*. Workflow redesign identified as the single strongest EBIT correlate of AI value capture — the empirical basis for treating the productivity uplift as capacity to be reallocated, not banked. <https://www.mckinsey.com/>

Carion et al., 2020. N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, S. Zagoruyko. *End-to-End Object Detection with Transformers (DETR)*. ECCV 2020. The set-prediction detection architecture adapted for borehole sinusoid geometry in this programme. <https://arxiv.org/abs/2005.12872>

Sculley et al., 2015. D. Sculley et al. *Hidden Technical Debt in Machine Learning Systems*. NeurIPS 2015. The canonical argument that the trained model is a small fraction of a production ML system — the reason ownership transfer dominates the long-run economics. <https://papers.nips.cc/paper/2015/hash/86df7dcfd896fcf2674f757a2463eba-Abstract.html>

From 3-4 Weeks to 2 Hours per Well: The Economic Case for Automating Borehole-Image Interpretation

Published April 2025. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/economic-case-automating-borehole-image-interpretation>

Copyright EarthScan. All rights reserved.