

Decommissioning at One-Third the Cost: A Simulation-Based Decision Framework

SEE THE PROCESS. CHANGE THE OUTCOME.

For offshore decommissioning directors, asset managers and procurement executives

Ilan Gleiser · Erica Crandon | Alpha IQ | August 2026

–69.4%

Well-work cost (P50)

–127 days

Campaign duration (P50)

–54.3%

Weather idle time

–73.7%

Scope 1 emissions

The bottom line: four decisions, made before any vessel sails, are worth \$71 million on a ten-well campaign

Static spreadsheets and linear Gantt schedules systematically misprice offshore plug-and-abandonment, because campaign outturns are driven by non-linear interactions — weather compounding with duration, crew learning compounding with batching — that deterministic tools cannot represent. Across 10,000 agent-based simulations per strategy of a ten-well subsea campaign, replacing the conventional programme (semisubmersible rig, sequential execution, section milling) with a simulation-optimized one (batched activity passes, light intervention vessel, perforate-wash-cement) reduces P50 well-work cost from \$102.5M to \$31.3M and completion time from 282 to 155 days.

These are ensemble medians from a validated model — twelve of twelve verification and validation checks passed, implementations bit-exact across two languages, comparative benchmark error of 3.2% — not point estimates from a plan. The distributions barely overlap: the optimized campaign's pessimistic case (P90) is still \$60M below the conventional campaign's optimistic case (P10).

Metric	Conventional (Rig + Milling)	Optimized (Batched LWIV + PWC)	Benefit
Total well-work cost (P50)	\$102.5M	\$31.3M	–\$71.2M (–69.4%)
Campaign duration (P50)	282 days	155 days	–127 days (–45.1%)
Waiting on weather (P50)	43.3 days	19.8 days	–23.5 days (–54.3%)
Fleet carbon emissions	16,760 tCO ₂ e	4,416 tCO ₂ e	–12,344 tCO₂e (–73.7%)

Four metrics, one decision — 10,000 simulated campaigns per strategy (P50, whiskers P10-P90)




Exhibit 1 — The two strategies barely overlap on any headline metric

Where the value comes from: a four-decision framework

The decision framework — four levers, one campaign

DECISION A	DECISION B	DECISION C	DECISION D
Campaign architecture Batched activity passes, not well-by-well execution	Vessel & technology Rigless LWIV + PWC vs semisub + section milling	Seasonal weather windowing Start dates and lift windows set against AR(1) metocean	Emissions & compliance Lighter tonnage, shorter campaign, less DP standby
MECHANISM Unbroken Crawford learning curves	MECHANISM USD 180k vs 350k per day; 2-5 vs 8-20 day baseline	MECHANISM Avoids compounding winter exposure	MECHANISM Fuel burn scales with engine-hours × tonnage
MEASURED ROI -45% duration	MEASURED ROI -69% cost	MEASURED ROI -54% weather downtime	MEASURED ROI -74% Scope 1 CO ₂ e

Source: welldecom agent-based simulation · 10-well subsea campaign · 12/12 validation checks · seeds-only bands

Exhibit 2 — Each decision maps to a mechanism the simulation quantifies

DECISION A — Batch by activity, not by well — and let crews ride one learning curve

The decision. Move from sequential execution (all phases on one well, then move the rig) to batched passes: reservoir isolation across all wells, then annular remediation, then conductor removal.

The simulation insight. Task-switching resets human learning. Batching creates an unbroken Crawford learning curve per crew — by the tenth repetition the same operation runs at roughly 60% of its first-well duration.

The measured ROI. 127 days removed from the campaign (282 → 155). Completion follows a plateau-then-sprint curve, with all wells reaching final abandonment in a rapid ~25-day succession — a signature to brief regulators on in advance.

DECISION B — Take the rig off the critical path: LWIV + PWC does the same barrier at half the day rate

The decision. Select light well-intervention vessels with single-trip perforate-wash-cement tooling instead of a semisubmersible with section milling.

The simulation insight. Milling takes 8–20 days per well with hot non-productive-time risks (swarf handling, stuck tools); PWC places the same annular barrier in 2–5 days with a colder risk table.

The measured ROI. Spread rate falls 48.6% (\$350k → \$180k per day) and the technology baseline shortens — together a \$71.2M net capital saving on the ten-well programme.

DECISION C — Pick the start date like a trade: weather exposure compounds with duration

The decision. Structure campaign start and lift windows against the seasonal, autoregressive metocean climate rather than a flat weather allowance.

The simulation insight. A campaign that slips samples more storms — and every waiting day is paid at the spread rate. Slipping into winter wave regimes (Hs > 2.5 m closes lift windows) triggers cascading delays a Gantt chart cannot see.

The measured ROI. Waiting-on-weather falls from 43.3 to 19.8 P50 days (–54.3%), preserving an estimated \$4.2–8.2M of paid vessel time; a winter start on identical scope roughly doubles weather downtime.

DECISION D — The same decisions that cut cost cut carbon — emissions are engine-hours

The decision. Optimize fleet tonnage, dynamic-positioning standby and campaign length against environmental and regulatory (e.g., IBAMA) commitments.

The simulation insight. Carbon intensity is driven by engine-hours times generator size. A shorter campaign on lighter tonnage compounds both terms at once.

The measured ROI. Scope 1 fleet emissions fall from 16,760 to 4,416 tCO₂e (–12,344 t, –73.7%), with time-on-location — the main driver of environmental exposure — down by more than half.

How to read these numbers — and what to do next

All figures are P50 medians from 10,000-iteration Monte Carlo ensembles with a fixed well inventory; whiskers in Exhibit 1 span P10–P90. The bands shown vary random seeds only — execution luck, not input assumptions. For budget-grade contingency, AlphaQ stacks plausible input ranges (day rates, learning rates, carbon price, metocean severity) on top, which widens the cost band by roughly 2×; every published band is labelled with what it varies, and every result reproduces from a recorded seed.

Recommended next steps

- 1.** Screen your well stock through the simulator: casing condition, water depths and target states drive which wells justify milling versus PWC.
- 2.** Run the strategy tournament on your portfolio — sequential vs batched, rig vs LWIV, start-date sweep — on ensembles, not single cases.
- 3.** Set contingency from the stacked (inputs + seeds) band at P90, and brief stakeholders on the phased completion signature before the first progress report.

Simulation results on stylised data, validated for model mechanics and comparative behaviour; deployment requires basin calibration. Not investment, provisioning or regulatory advice. © AlphaQ 2026.