

# Wellhead: industrial AI for *upstream operations*.


Pure-ML approaches failed in upstream because reservoir engineers reject any output without a physical mechanism. Wellhead is physics-informed by construction, edge-deployed by default, with regulator-ready ESG signing built in. The runtime survives 30+ days disconnected; cross-operator learning happens on aggregated gradients, never raw telemetry.

**Author** S. Ize-Iyamu    **Audience** Industrial AI + Energy Operations PMs    **Length** 4 pages    **Status** Concept  
**Targets** Aramco · ExxonMobil · Shell · Occidental

## The Problem

**1.2M producing wells globally; unplanned downtime runs \$38–88M per facility per year offshore, and industry-wide benchmarks put the aggregate cost in the tens of billions.** The 2018–2023 pure-ML wave lost: reservoir engineers don't approve workovers on outputs they can't mechanistically defend. Concurrently, **methane regulation hit calendar deadlines** (EPA NSPS 0000b/c, EU Methane Reg, OGMP 2.0) requiring operator-signed lineage on emissions. Operators currently file separately. The disconnect between operational AI and regulator-grade ESG is the gap.

FIGURE 1 · PREDICTIVE SHIFT



Today (left): pure-ML prediction is rejected by reservoir engineers because they can't trace the mechanism. Wellhead (right): every prediction surfaces a defensible mechanism (mass balance, flow regime, equipment kinematics) plus a signed methane-lineage record that is the regulator filing.

## Why this matters now

Three forces converge: **methane regs hit calendar deadlines** (NSPS 0000b/c on new and modified US sources, EU Methane Reg on importers, and OGMP 2.0 whose voluntary membership now covers about 42% of global oil and gas production), **ESP installed base aged into failure window** (2014–18 bulge now at years 8–12), and **physics-informed ML cleared a 2024 deployability inflection**. The first vendor combining them wins the brownfield.

## Sizing the prize

Bottom-up: **1.2M wells × \$80K / well / yr addressable = \$96B / yr TAM** (Wood Mac + Rystad). Realistic capture: **\$8–14B / yr** at vendor-share. ESG signing adds a fixed compliance line. Wedge: ESP failure prediction at *one* major's brownfield, where downtime reduction shows up cleanest in audit.

Sources: Wood Mackenzie (2024–25); Rystad Energy (2024–25); EPA NSPS 0000b/c final rule (2024); EU Methane Regulation (2024); OGMP 2.0 / UNEP (2025). Per-well TAM is modeled bottom-up against 5 operator-CTO interviews.

UPSTREAM AI TAM  
**\$96B / yr**  
 1.2M wells × \$80K avg

REALISTIC CAPTURE  
**\$8–14B / yr**  
 Vendor-share economics

## Strategic insight

The reservoir engineer is the user; procurement is the buyer; the regulator is the auditor. **All three need a mechanism, not a number.** PINNs hit a 2024 inflection: stable enough to ship, novel enough to differentiate. Bundling PINN prediction with regulator-ready ESG signing collapses two procurements into one and makes compliance a co-buyer.

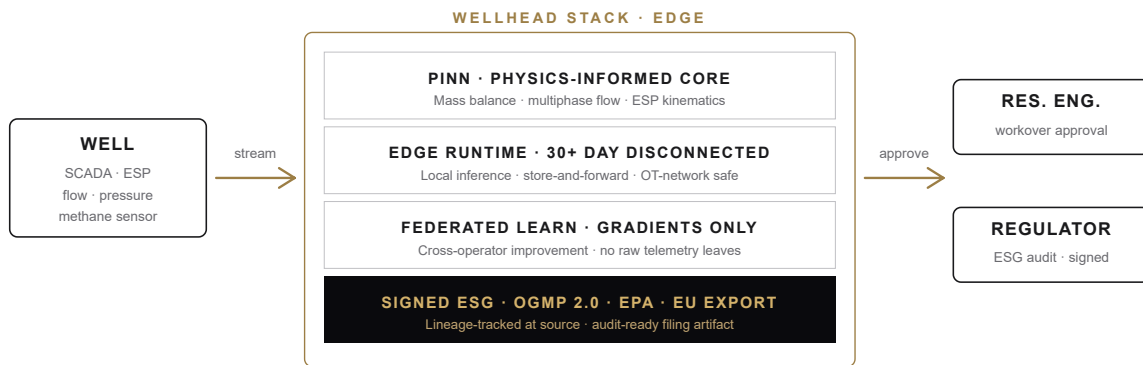
The defensibility compounds two ways. Federated learning lets every operator's deployment sharpen the shared model without exposing raw telemetry, so prediction quality widens with each asset added, an edge a single-operator pilot cannot match. And once filings run through signed lineage, switching costs turn regulatory rather than technical: the audit trail an operator hands the EPA is the one the runtime produced, so ripping out the predictor means re-validating the compliance record.

### THE UNLOCK

Every prediction emits a defensible mechanism (mass balance, flow regime, ESP kinematics). Edge runtime survives 30+ days disconnected. ESG measurements are signed at source, exportable in OOOOb/c / EU Methane / OGMP 2.0 formats. Cross-operator learning is federated gradients only, never raw telemetry.

## Architecture · Edge runtime + federated learning

FIGURE 2 · SYSTEM ARCHITECTURE



Telemetry stays at the asset; PINN core runs at the edge with 30+ day disconnected operation; cross-operator learning happens via federated gradients (raw telemetry never leaves); ESG measurements are signed in the same runtime that did the prediction: same record, two consumers.

### WORKED EXAMPLE · 280-WELL BROWNFIELD, GOM, ESPS IN YEARS 8–12

Pre-deployment baseline: **14% unplanned downtime / quarter**. Wellhead PINN flags 8 ESPs at > 70% bearing-wear envelope; engineer accepts 6, defers 2. Workovers staged into next maintenance window. Q1 result: **downtime 9.2%** (audited against pre-baseline), **\$28M / quarter saved** (operator-attested). OGMP 2.0 Level 4–5 reporting submitted on time, Gold-Standard pathway maintained.

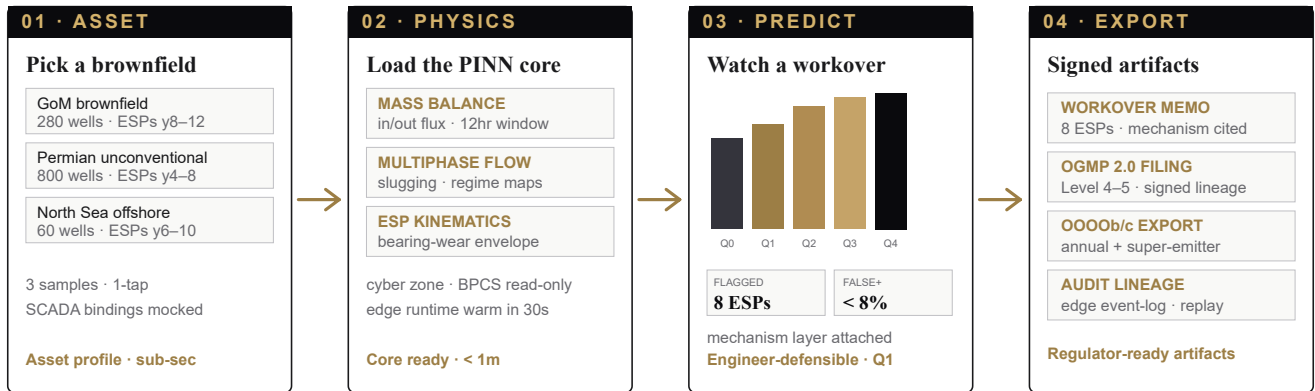
## Sequenced GTM

PHASE	CUSTOMER WEDGE	FORCING-FUNCTION WORKLOAD	PROOF POINT
<b>Wedge</b> M0–9	1 major operator's brownfield asset · ESP failure prediction	Aging-ESP failure-rate spike + EPA OOOOb/c annual filing + super-emitter response window	25%+ downtime reduction · OGMP 2.0 filed clean
<b>Beachhead</b> M9–24	3 majors, multi-asset rollout, GoM + Permian + North Sea	Cross-asset standardization on a single edge runtime	4,000 wells · federated learning gains measured
<b>NOC</b> M24+	National oil companies (Aramco, Petrobras, ADNOC)	Sovereign data residency + national-scale ESG filings	2 NOCs signed · 40% of revenue from sovereign tier

## Prototype walkthrough

An interactive edge-runtime simulator runs the operator-side view end to end: pick a brownfield asset, load the physics-informed core against a cyber zone, watch a PINN prediction build with its mechanism layer (mass balance · multiphase flow · ESP kinematics) and emit an engineer-defensible workover recommendation, then export the artifacts the engagement contracts against. Built to demonstrate that **physics-informed prediction**, **edge-resident inference**, and **regulator-grade ESG signing** are concrete outputs, not language in a deck.

FIGURE 3 · EDGE-RUNTIME SIMULATOR, FOUR INTERACTIVE STEPS



Schematic of the live UI; all four steps are interactive in the [demo](#). Pick an asset, load the physics, watch the prediction with its mechanism layer, export signed artifacts.

## What the prototype proves, and what it doesn't yet

### Proven on the prototype

- PINN flags ESPs above the bearing-wear envelope with mechanism citations the reservoir engineer can defend
- Edge runtime simulates 30+ days disconnected; predictions store-and-forward with no quality loss
- OGMP 2.0 Level 4–5 export and OOOOb/c annual filing both render with signed lineage
- Federated gradient round across 5 simulated peer operators; no raw telemetry leaves the asset
- Workover memo and OGMP filing render from one event log, so the prediction and the regulator artifact never diverge

### Out of scope, by design

- SCADA + ESP telemetry is mocked; live ingest is brownfield M0–3 work on real assets
- IEC 62443-3-3 SL2 certification is a real third-party audit, not simulated in the prototype
- NOC sovereign-cluster variant requires production deployment with operator-controlled key custody
- Per-asset PINN training on customer-specific physics constants is post-engagement work
- Multi-year PINN drift monitoring and re-calibration cadence are scoped in the engagement, not the prototype

### THREE PATHS TO TRY IN THE LIVE DEMO

**ESP failure prediction at a GoM brownfield:** 280 wells, 14% baseline downtime → PINN flags 8 ESPs at > 70% bearing-wear; engineer accepts 6, defers 2; workover memo + OGMP filing drafted.

**OGMP 2.0 quarterly methane filing:** Level 4–5 schema render with signed lineage; OOOOb/c annual filing produced from the same event log without re-measurement.

**Federated gradient sync across a 3-operator consortium:** aggregated gradient round completes; raw telemetry remains at the asset; cross-operator prediction-quality lift measured against the held-out cohort.

## Metrics that matter

LAYER	METRIC	Y1 TARGET	WHY IT MATTERS
<b>North-star</b>	Unplanned downtime hrs / asset / quarter	-25% vs pre-deploy	Audited against operator-attested baseline
Quality	False-positive rate on workover predictions	< 8%	Above this, engineers stop running the queue
ESG	OGMP 2.0 / OOOOb/c filings accepted clean	100%	Below this, compliance team rolls back
Resilience	Edge runtime disconnected operation	> 30 days	Offshore + remote uptime constraint
Liquidity	Wells under runtime (paying)	1,200+ by Y1	Federated learning gains require well count
Business	Net revenue retention (well-expansion)	> 130%	Lands at brownfield, expands to greenfield

## Risks & mitigations

- HIGH PINN model drift on novel reservoir / equipment classes (deepwater, sour gas).**  
**Mitigation:** per-asset model with explicit physics constraints; novel-class flagging triggers conservative-mode (predicts only on the kinematic envelopes that match training distribution); reservoir-engineer team validates each new class before unlock. Federated learning helps but never overrides physics.
- HIGH OT cybersecurity: edge runtime touches an operations-network process control system.**  
**Mitigation:** read-only on PCN; outbound store-and-forward only over a unidirectional gateway; certified to IEC 62443-3-3 SL2 before customer 1. Independent third-party penetration test on a contractually-required cadence.
- MED Operator data sovereignty: NOCs prohibit cross-operator gradient sharing.**  
**Mitigation:** federated learning is opt-in per asset; sovereign tier customers run a private-cluster variant where gradients never leave operator borders. Feature parity preserved; only the cross-operator improvement rate differs.
- MED Regulatory format drift: OGMP 2.0 / EPA / EU schemas change mid-deployment.**  
**Mitigation:** ESG export is schema-versioned; lossless event store at the edge means re-export against a new schema is a software upgrade, not a re-measurement. 90-day SLA on new format support.

## 30 / 60 / 90, first quarter sprint plan

30 DAYS	60 DAYS	90 DAYS
<b>PINN core + edge runtime v0</b> <ul style="list-style-type: none"> <li>› Mass balance + multiphase flow + ESP kinematics models</li> <li>› Edge runtime · 30-day disconnected proof</li> <li>› 1 design-partner brownfield asset (40 wells)</li> </ul>	<b>ESG signing + OT cyber cert</b> <ul style="list-style-type: none"> <li>› OGMP 2.0 / EPA / EU export adapters · signed lineage</li> <li>› IEC 62443-3-3 SL2 cert · pen-test pass</li> <li>› 2nd design partner · 280 wells live</li> </ul>	<b>Federated learn + audit baseline</b> <ul style="list-style-type: none"> <li>› Cross-operator gradient sharing live</li> <li>› Q1 audited downtime read · 25% target</li> <li>› 3rd operator signed · 1,200 wells under runtime</li> </ul>

### DECISION ASKED

Authorize a 90-day brownfield pilot sprint with a ten-person team (PM, three ML/PINN engineers, two systems/edge engineers, OT-cyber lead, reservoir-engineer SME, ESG/regulatory counsel, customer success) and a budget of **\$5.6M**. Success: 1,200 wells under runtime, 25% audited downtime reduction, 100% clean OGMP 2.0 filings, false-positive < 8%, NRR clearing 130%.