

Where LLMs can help create value.

LLM contribution

Potential
economic
impact

Small supporting role

Useful for part of the work

Central to much of the task

High

01 Lift optimization
ExxonMobil · Devon · Chord

02 Plant control
Aramco / Fadhili

04 Equipment reliability
Equinor · Murphy Oil

03 Field planning
Equinor · Petrobras · Quantum

05 Production exceptions
Presidio · Occidental

Medium

06 Procurement / rebids
EQT / prospective opportunity

07 Financial recovery
SLB / export duties

08 Data validation
Continental Resources

09 Revenue accounting
Spearpoint Resources

Lower

10 Business drafting
PRIO / HR documents

Role of an LLM in the workflow

■ Keep the specialist engine

■ Extend the operating workflow

■ Language-heavy work

QUALITATIVE ANALYST ASSESSMENT / Ten selected workflow families. These are not measured returns or precise scores.

Company cases support the existing workflows; proposed LLM roles require testing against the current system. Scale can change a placement.

Where LLMs can help create value.

Agent with tools: proposed role with persistent cases and approved tools

Potential
economic
impact

Small supporting role

Useful for part of the work

Central to much of the task

High

02 Plant control
Aramco / Fadhili

01 Lift optimization
ExxonMobil · Devon · Chord

05 Production exceptions
Presidio · Occidental

03 Field planning
Equinor · Petrobras · Quantum

04 Equipment reliability
Equinor · Murphy Oil

Medium

06 Procurement / rebids
EQT / prospective opportunity

07 Financial recovery
SLB / export duties

08 Data validation
Continental Resources

09 Revenue accounting
Spearpoint Resources

Lower

10 Business drafting
PRIO / HR documents

Role of an agent with tools

■ Keep the specialist engine

■ Extend the operating workflow

■ Language-heavy work

QUALITATIVE ANALYST ASSESSMENT / Ten selected workflow families. These are not measured returns or precise scores.

Company cases support the existing workflows; proposed LLM roles require testing against the current system. Scale can change a placement.

Two judgments. No ROI score.

Placement basis

Qualitative analyst assessment of ten selected workflow families, informed by the linked public-source cases. Company examples establish disclosed workflows and reported outcomes, not the placements or the incremental return of a proposed LLM redesign.

Economic impact

Potential economic impact for a suitably scoped operator deployment: High means material production, availability or capital decisions; Medium means recurring financial, commercial or information work; Lower means an individual content task. These bands are contextual hypotheses, not measured ROI, market size or universal rankings. A large financial workflow can exceed a small operating deployment.

LLM fit

Small means a supporting role with the specialist engine central; Part means useful interpretation or investigation while other systems complete the job; Central means language and document interpretation are central to the selected task. Tools do not turn an LLM into a validated numerical optimizer or equipment controller.

Agent view

The agent view assumes permissioned access to records and specialist tools, persistent case state, event-driven resumption, checks and human approval where needed. It is a proposed architecture, not proof of autonomous field deployment or higher returns.

The CSV and JSON downloads include all placements, reasons, reported metrics and source boundaries.

01 / Lift optimization

A small improvement that lasts across many suitable wells can add production and reduce operating costs. The main job is to understand the well and calculate the right settings. A language model can help people investigate problems and arrange a response.

ExxonMobil / AI-235: Operator-authored SPE abstract reports about 2% average oil uplift. JPT's October 2024 interview reports 1.87% in a three-month human-approval stage and about 2.2% in the next three-month, mostly automatic stage. These summaries are not additive; periods and cohorts are incompletely disclosed. Boundary: The primary paper's accessible abstract states the full paper is unavailable, so full technical validation is not possible. The operator-reported uplift lacks a disclosed counterfactual or confidence interval. Gas-lift performance can drift when water cut or productivity changes; JPT says an existing setting may remain suboptimal for years without frequent monitoring. The control-boundary details and automatic-versus-human exception rates are undisclosed.

Original disclosure / 2025-10-01

Devon Energy / AI-006: 1,000 live wells as of August 2026; >2,000 additional wells prospective. Separate 2025 gas-lift pilot reported 2-3% uplift in May 2026; not a measured result for the full August fleet. Boundary: Transcript reports management claims; no eligible-well denominator, accuracy, lift mix by well, independent validation, or AI-specific production delta disclosed. Broader >2,000 is opportunity, not deployed scale. Do not attribute the company-wide business optimization program to AI.

Original disclosure / 2026-08-05

Chord Energy / AI-097: Operator reports early approximately 25% improvement in rod-pump run times. The 99% denominator is rod-lift wells, not the entire well fleet. Boundary: Company-reported early result without a disclosed comparison method. Broader workover, ESP and savings results cannot be isolated to AI.

Original disclosure / 2026-02-25

02 / Plant process control

A plant uses steam, chemicals and electricity every day. Reducing that use can make a meaningful difference to its costs. The control system has to keep the physical process running reliably. A language model has not been shown to choose better settings.

Aramco / AI-141: amine and steam use reduction: 10%-15% (Company-reported); power use reduction: around 5% (Company-reported) Boundary: Metric attribution is company-reported and evaluation details are not disclosed. The company source page is undated; accessed 2026-09-27. The specific month permanent adoption occurred is not disclosed.

Original disclosure / Undated; accessed 2026-09-27

03 / Well and field planning

A different development plan can change how much capital is spent, how much is produced and when the money comes back. A language model can help explain assumptions and compare results. Engineers and specialist software still need to check that a plan can work.

Equinor / AI-035: \$12m partnership savings Boundary: Included case study, not incremental to portfolio total.

Original disclosure / 2026-01-07

Petrobras / AI-219: R\$350 million captured; period and definition of captured value unstated. Boundary: No baseline, project duration, well count or value-calculation method. Separate from Petrobras flare control and monitoring.

Original disclosure / 2025-11

Quantum Capital Group / AI-054: Microsoft reports a 20,000-acre scenario took 20 minutes versus three weeks; leasehold utilization rose from 50% to 85%. Production runs evaluated 2.3m pad candidates, 3.5m well possibilities and over 50m allocations; the accumulation period is unstated. Boundary: The operator is unnamed. Optimization and a conversational interface are separate components; no autonomous capital approval is shown. Vendor-reported comparisons lack independent validation, exact coverage and full costs. Parallel optimization of multiple scenarios is prospective.

Original disclosure / 2026-04-30

04 / Equipment reliability

Avoiding a serious outage can protect production and reduce repair costs. Finding early signs of equipment failure usually needs a specialist model. Reading documents alone does not reliably predict a failure.

Equinor / AI-034: >700 machines; \$120m value since 2020 Boundary: Cumulative; overlaps overall value.

Original disclosure / 2026-01-07

Murphy Oil: The first model took six months to deploy because required data was missing. The case gives no before-and-after false-alert rate or measured net return.

JPT / SPE: offshore predictive-maintenance case

05 / Production exception response

Fixing recurring problems across a group of wells can recover production and change how field crews spend their time. A language model can bring the relevant information together. Someone still has to act on its recommendation and check what happened.

Presidio Production / AI-231: Operator reports 2.3% AI-linked production uplift through Q2: 1.4 points from DOUG and nearly 1 point from adjacent initiatives. DOUG averaged ~400 boe/d incremental production over the prior three months. \$4.5m annualized revenue spans the initiatives; it is not net cash or a DOUG-only result. Boundary: Company-reported attribution; no independent counterfactual, measurement protocol or project costs. The \$4.5m annualized revenue spans AI initiatives and is not realized net cash.

Original disclosure / 2026-08-12

Occidental / AI-215: Approximately 40% of onshore production covered; this is a production-share denominator, not a share of wells or personnel. Boundary: No route-level time, cost, failure-rate or production-uplift result. The contribution of AI cannot be separated from sensors and automation.

Original disclosure / 2025-08-07

06 / Procurement and rebids

Choosing better contracts can affect substantial spending. Savings still need to preserve the scope and quality of the service. A language model can help read the documents. Savings depend on comparing the offers fairly, making a decision and completing the change.

EQT / AI-095: Prospective low-single-digit procurement cost reduction; unit and spend denominator unspecified Boundary: Low-single-digit procurement cost reduction is an opportunity estimate, not a realized saving or an isolated AI effect. Tool, vendor, scope and baseline remain undisclosed.

Original disclosure / 2026-02-18

07 / Document-led financial recovery

A well-supported claim can recover money. How much depends on the number of transactions and which ones qualify. Reading the documents is central to this job. The records still need to be reconciled using checked rules and reviewed by finance.

SLB / AI-234: Microsoft page reports \$10 million recovered in export duties; workflow description says it processes received documents. The source does not disclose the measurement period, baseline, or direct attribution method. Boundary: Microsoft-hosted customer story, not independently validated. The \$10 million figure has no stated period, reconciliation, or counterfactual. The exact publication date is not explicit beyond page last-updated 2025-08-08; no launch date or throughput/accuracy is disclosed. This is a distinct shipping-document/export-duty workflow, not the airline invoice GST bot also described on the same page.

Original disclosure / 2025-08-08

08 / Operational data validation

Less manual checking could free up engineers' time and reduce delays or work that has to be redone later. This job depends on understanding documents and language that do not always arrive in the same format.

Continental Resources / AI-248: Undated vendor case: 200+ emails/day, approximately five minutes saved per previously manual email, one reusable robot supporting 40+ processes. Undated award page: accuracy rose from about 70% to the high 90% range. Boundary: 2025-09-29 dates the award announcement only. Current workflow and metric dates are unknown; 40+ processes are not 40 AI initiatives. No test-set definition, error monitoring or independently verified cash savings.

Original disclosure / Undated case; award announced 2025-09-29

09 / Revenue-statement reconciliation

This work happens repeatedly. Doing it better could free up time, help close the books sooner and find payment discrepancies. Reading the statements is central to the job. Arithmetic and accounting judgments still need their own checks.

Spearpoint Resources / AI-238: Spearpoint Controller Elle Ross estimates Joltly gives the team 20–30 hours per week back across three entities. From May through August 2026 the case reports 82 statements processed, 5,441 pages and 56,358 line items across 27 purchaser formats. Time estimate is customer-reported; no baseline time study, quality/exception rate or independent validation is supplied. Boundary: Vendor-authored case study; no audited results, accuracy/error rate, control testing, implementation cost, customer-specific pricing, or independent attribution. The output is reviewed/exported by people; do not describe it as fully autonomous posting. It is a named but small non-operated asset owner described without basin or asset-level details.

Original disclosure / 2026-09-05

10 / General business drafting

Here we are looking at one writing task. Benefits across a company could be larger, but time saved does not automatically reduce cash costs. Writing is the main job, so a language model can do much of it without an industrial optimizer.

PRIO / AI-277: PRIO staff report more than 70 hours saved per area: approximately 50 consultant hours plus 26 manager hours previously, versus about four hours for the new process. A claimed 99% “assertiveness” has no stated validation method. Boundary: Vendor-published company quotes; approximate baseline. Number of areas, aggregate hours, quality-test denominator, subscription costs and realized cash savings are not supplied. Historical-origin implementation predates the reporting window.

Original disclosure / 2025-04-15