McKinsey & Company

Oil & Gas Practice

Can upstream oil and gas produce value from AI's $230 billion pay zone?

AI could unlock $230 billion for oil and gas and OFSE companies. The challenge is where to concentrate, how to scale, and how to share value when efficiency reduces the activity that many contracts reward.

by Bill Ambrose, Giorgio Bresciani, and Spandan with Priyank Singh

August 2026

For much of the past decade, digital transformation in upstream oil and gas has delivered real but uneven results. Efficiency has improved, operations are safer, and data infrastructure is more mature. Yet many executives at operator and oilfield services and equipment (OFSE) companies are asking: Where is the step change in value?[1]

AI is beginning to provide an answer although its value is not uniformly distributed. A bottom-up sizing of AI use cases across value chain stages estimates that AI can unlock approximately $65 billion in annual recurring value across upstream oil and gas in the near term with today's technology, and there is a credible path to $230 billion at full potential as the technology matures and autonomous operating modes become more common.[2] In addition, AI-driven improvements in exploration success could unlock over $35 billion annually in balance sheet value through reserves accretion.

This estimate is built from more than 550 use cases across key domains in the upstream life cycle (see sidebar, "How we sized the opportunity and defined the scenarios"). Yet, AI is not free. These estimates are net of annual implementation costs that include compute, specialized talent, data infrastructure, and software, totaling over $30 billion.

AI in upstream is a concentration play, not a "thousand flowers bloom" opportunity. The top ten use cases drive nearly half of the identified value, the top 20 drive roughly two-thirds, and the top 60 capture approximately 95 percent of identified value (Exhibit 1).

[1] Bill Ambrose, Spandan, Priyank Singh, and Sri Kandala, "Gen AI in the OFSE industry: Progress lags behind intent," McKinsey, April 1, 2026.

[2] Full potential assumes autonomous operating modes and maximum addressability across all life cycle stages within the upstream oil and gas segment. A cascade correction is applied from gross stand-alone benefits of each use case to avoid double counting use cases acting on the same wells and operations. These estimates are net of annual implementation cost, including compute, specialized talent, data infrastructure, and domain software of over $30 billion.

Exhibit 1

The top ten AI use cases drive nearly half of the identified value in upstream oil and gas.

Value concentration of AI use cases, %

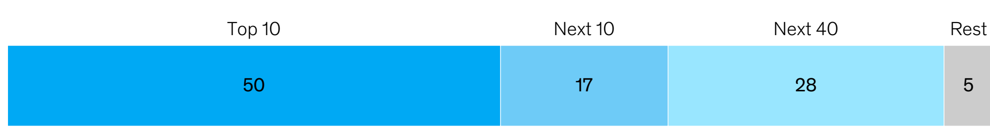

Note: McKinsey analysis of 550 use cases.

McKinsey & Company

How we sized the opportunity and defined the scenarios

AI can drive value for the upstream oil and gas sector in four main ways: increasing production, reducing operating cost, improving capital productivity, and changing the risk and balance sheet profile of assets. Our analysis defines value through the operator's revenues, expenses, capital expenditure, and balance sheet impact including increased production from existing wells, lower nonproductive time, lower lease operating expenses, fewer wells required to meet development plans, improved recovery, lower inventory, and shorter field development cycles.

Because operators rely on oilfield services and equipment (OFSE) companies to deliver against their plans, we also sized the opportunity by OFSE subsegment. For oilfield services, AI products that create the most value are typically concentrated in drilling, production optimization, and maintenance—where data infrastructure and closed-loop feedback already exist.

Our estimates are calculated bottom-up from over 550 individual AI use cases mapped across the upstream life cycle and applied to the oil and gas industry's global cost and production base:

- cost base: approximately $500 billion to $550 billion in annual capital expenditure and approximately $400 billion to $500 billion in annual operating expenditure
- production base: approximately $2 trillion in annual production margin at $50 per barrel of oil equivalent (BOE)[1]

Each use case was sized by estimating the fraction of total spend or production value that AI can act upon today, the realistic deployment rate, and the estimated improvement range. Estimates were anchored to operator and service company proof points and corrected for overlap across use cases.

The use cases were sized with named proof points.[2] A cascade correction prevents double counting where multiple use cases act on the same value pools.

The near-term potential of approximately $65 billion reflects today's proven AI technologies, applied at current adoption levels.

The medium-term potential of approximately $125 billion per year still includes proven technologies that are adopted more broadly as platform economics improve and early-adopter results drive wider deployment. Addressability approximately doubles as connectivity, data standards, and workforce readiness improve.

The full-term potential of approximately $230 billion assumes autonomous operating modes and maximum addressability across the value chain. This scenario considers closed-loop control where physics permits and commercial models that reward scaled deployment (exhibit).

---

[1] Global upstream capital and operating expenditure data 2024–2025, Rystad Energy.

[2] The use cases cover multiple AI categories, including machine learning (ML: classification, regression, optimization on structured operational data; physics-informed ML (models constrained by engineering equations so that outputs remain physically defensible; deep learning and surrogate models (neural networks for signal interpretation, image analysis, and fast approximation of physics simulators; computer vision (image and video analysis for inspection, core analysis, seismic interpretation, and remote monitoring; natural language processing (extraction, classification, and search over unstructured reports and documents; gen AI (large language models for knowledge synthesis, report generation, and scenario creation; and agentic AI (systems that autonomously plan and execute multistep operational workflows.

Exhibit

The gap between ‘near-term potential’ and ‘full-term potential’ closes as adoption broadens and technology matures.

Gap between near-term potential and full-term potential, $ billion per year

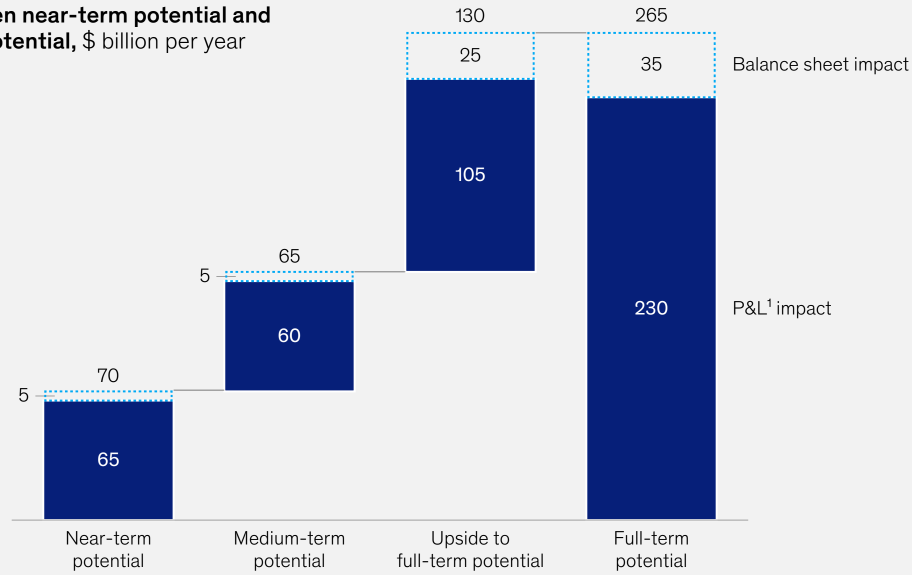

[1]Profit and loss.

McKinsey & Company

The implication is that the companies that are likely to succeed in capturing value from AI will be those that focus relentlessly on the handful of use cases where potential value is greatest—those with large value pools, real-time data, physics-based feedback, and measurable operating KPIs. McKinsey analysis shows a steep drop-off in value opportunity beyond the top use cases (Exhibit 2).

Exhibit 2

AI in upstream is a concentration play, as most value is trapped in the top use cases.

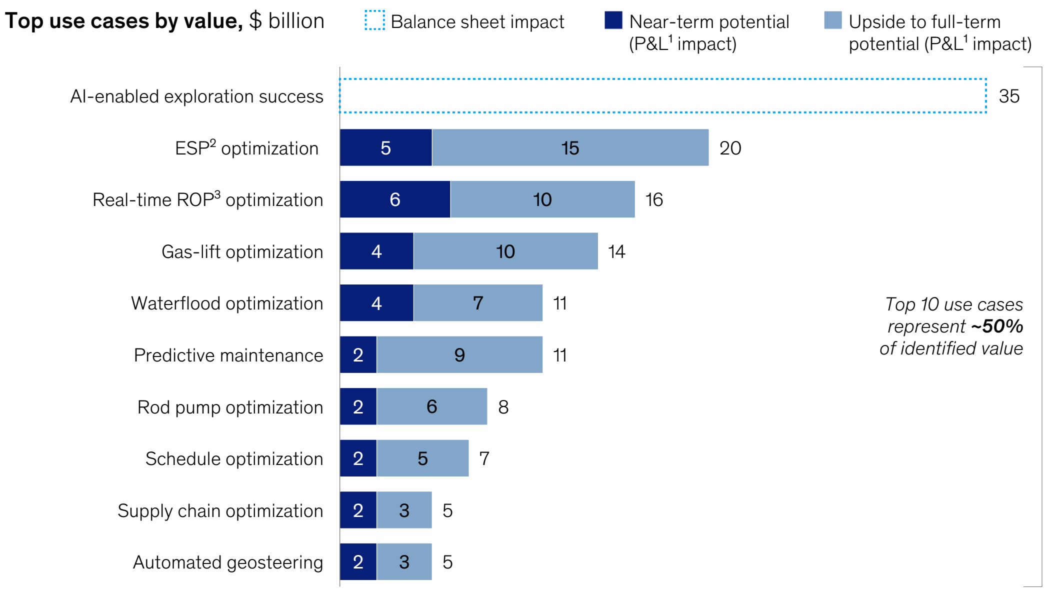

[1]Profit and loss.
[2]Electric submersible pumps.
[3]Rate of penetration.

McKinsey & Company

The value is concentrated in the operating core

Value tends to concentrate in the develop and produce life cycle stages, with the find stage benefiting from balance sheet impact through incremental reserve accretion (Exhibit 3).

Simply put, our analysis shows that the largest near-term prize sits in the develop and produce stages of the value chain. These stages share common features. They represent large pools of economic value, they generate high-frequency operational data, and they are governed by physical systems where feedback can be observed. Plus, they have measurable metrics, including cycle time, uptime, recovery, cost per foot, cost per barrel, nonproductive time (NPT), and capital efficiency.

The top use cases fall into three operational domains—production optimization, drilling activity, and reservoir management—that together account for most of the value. Understanding these domains, rather than individual use cases in isolation, is the key to capturing AI value at scale.

Autonomous artificial lift and production optimization. Production assets generate highfrequency operational data with immediate physical feedback, which makes them ideal terrain for AI to add value. AI can optimize rod pumps, electric submersible pumps (ESPs), gas lift, waterfloods, production networks, surface facilities, flow assurance, and chemical programs. Value from these use cases can be measured quickly and effectively through incremental production, lower well downtime, fewer failures, lower intervention costs, and lower energy intensity.

Exhibit 3

Three stages of the value chain capture the bulk of AI value, which means that prioritization, not breadth, determines return on investment.

Value by life cycle stage, near-term potential, $ billion

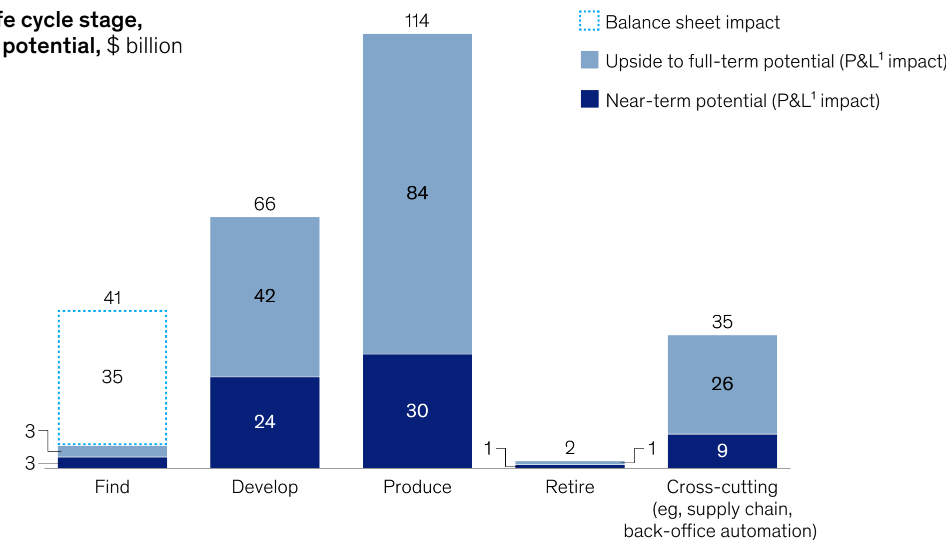

[1]Profit and loss.

McKinsey & Company

Drilling and well delivery. Drilling remains one of upstream's largest outsourced cost pools, with significant spend in offshore and land contract drilling, directional drilling, fluids, cementing, bits, and logging.[3] AI can improve well planning and equipment selection, reduce tripping and connection time, provide real-time and early anomaly detection, and execute closed-loop drilling automation for optimal well placement within the reservoir. Together, these applications can shorten the execution timeline through efficient operational sequencing, reduced nonproductive time, and optimal productive time management. The result is a reduced cost per day to deliver a well.

Intelligent reservoir management. AI can accelerate subsurface interpretation, improve static and dynamic model updates, generate surrogate models for faster simulation, support recovery strategy, improve reserves estimation, and help teams make better decisions. In reservoir management, machine learning (ML), ML-plus-physics models, deep-learning surrogate models, gen AI, and agentic AI systems can plan and execute multistep workflows with limited human intervention, which can reduce time taken for exploration and field development from months and years to days and weeks. But this opportunity for AI applications will require significant change in how activities occur today across reservoir planning and management that are grounded in highly siloed workflows and experience across organizations.

[3] McKinsey analysis.

The industry is already seeing value from AI deployment. For example, SLB and Vår Energi arepooling resources to build collaborative well-planning workflows that reduce cycle times from discovery to first oil from months to days for projects on the Norwegian Continental Shelf.[4] Baker Hughes and Expand Energy announced a multiyear collaboration to deploy Leucipa™, an AI-powered automated field production solution, across thousands of natural gas wells, piloting an AI production assistant to enhance operational agility.[5]

Is the industry solving the wrong scaling challenge?

The industry's AI challenge is often described as a technology problem: incomplete data, legacy systems, limited connectivity, and models that are not robust enough for complex field environments. These barriers are real, but many companies may be tackling the wrong scaling problem. AI creates the most value when it changes decisions. The scaling challenge is to build new ways of working that use AI tools to make faster, more consistent, and higher quality decisions that drive better outcomes with lower risk profiles.

In our work, we have seen several factors that hinder upstream companies from scaling AI:

— **Prioritization is too broad.** Many portfolios still include dozens of digital and AI initiatives, often sponsored by different functions and justified by local productivity gains. But given that value is a concentration game, spreading investment across the long tail is a form of value dilution.

— **Build-versus-buy decisions are often driven by control.** Operators and OFSE companies frequently build technology capabilities internally, such as data platforms, surveillance workflows, drilling analytics, and maintenance models, because they believe their assets or operating contexts are unique. In certain cases, that is true. But excessive internal build can be expensive and slow, and makes it harder to take advantage of broader, industry-wide data sets that may be required to deliver accuracy of AI models in operational contexts.

— **Change management is underestimated.** AI use cases that require frontline adoption, such as drilling parameter optimization, maintenance planning, and production surveillance, depend on trust and workflow redesign. A model can be technically correct yet still fail if the user does not trust the answer or does not change their ways of working and making decisions.

— **Value tracking can be weak.** McKinsey analysis indicates that less than 20 percent of deployments track AI-generated value and share the KPI with their organizations. Without transparent baselines and agreed-upon methods to measure value, AI programs remain vulnerable to skepticism.

— **Finally, and most important, the commercial model can be misaligned.** Under many industry operator service provider contracts, AI creates a paradox. Operators benefit from faster wells, fewer failures, lower downtime, and fewer interventions, but OFSEs may lose

---

[4] SLB and Vår Energi expand digital collaboration to scale well and integrated field development planning," SLB press release, May 28, 2026.
[5] Baker Hughes, Expand Energy to deploy Leucipa across US natural gas fields," Baker Hughes, January 29, 2026.

revenue when the same efficiency reduces their billable rig days, service hours, frac stages, maintenance events, or time-and-materials activity. A commercial model that shares the value upside can help to incentivize the use of efficiency-creating AI.

From the OFSE perspective, AI is a double-edged sword. On the one hand, at current AI deployment levels, approximately $17 billion of oilfield services' revenue could be exposed as AI-driven efficiency reduces activity under existing commercial structures.[6] At full potential, that revenue exposure could reach $60 billion.[7] However, revenue loss overstates the actual financial impact. When activity declines, OFSEs also shed the variable costs associated with it—crew, equipment, consumables, and logistics. McKinsey analysis shows that at typical industry margins of about 40 percent, the real cash flow impact is closer to $7 billion at near-term potential and $24 billion at full potential.

OFSEs that use AI to enhance their operations could improve their efficiency by approximately $10 billion today—and up to $20 billion at full potential. This largely closes the cash flow gap and points to the real strategic question: How can AI help make OFSEs better off than they are today?

OFSEs that develop and deploy technology are not always rewarded for the corresponding efficiency. To illustrate the problem, consider drilling. If closed-loop drilling automation reduces days on well, the operator benefits immediately. But under a traditional day-rate or activitybased model, the driller and service companies earn less—unless the contractual relationship shares the gain. Also, in production maintenance, predictive analytics can reduce failures and field activity, but fewer well interventions can mean fewer billable events for OFSE providers. In completions, better designs and execution can reduce rework or cycle time, but the benefit may accrue disproportionately to the operator rather than to the OFSE provider that creates the AI solution.

Production uplift may be the exception. Incremental barrels create new value without necessarily displacing OFSE revenue. That is why production optimization, artificial lift, and reservoir management are natural candidates for performance-linked models. Our analysis identifies a production uplift opportunity of up to $35 billion and notes that incremental barrels are more naturally aligned to drive shared performance outcome arrangements than efficiency savings that reduce billable activity.[8]

Capturing AI's untapped potential

Capturing the next wave of AI value in upstream requires a change in approach. The industry may not need more pilots or proofs of concept. It likely needs fewer and bigger bets, clearer ownership of solutions and data, industrialized deployment, and commercial models that reward outcomes—and share the value. As McKinsey's *Rewired* framework argues, technology-enabled transformation works when strategy, talent, operating model, technology, data, and adoption and scaling are built together.[9] Each of these areas has implications for companies in oil and gas.

---

[6] McKinsey analysis, based on oilfield services' share of AI-addressable activity, calculated cluster by cluster across drilling, completions, production, equipment, and maintenance service lines. Analysis was conducted using more than 10,000 statements for a range of OFSE companies.

[7] McKinsey analysis.

[8] Production uplift use cases totaling $15 billion to $34 billion; incremental barrels create new revenue without displacing existing service activity; McKinsey analysis.

[9] Eric Lamarre, Kate Smaje, and Rodney Zemmel, *Rewired: The McKinsey Guide to Outcompeting in the Age of Digital and AI*, June 2023.

Strategy: Concentrate on the few domains that matter

Operators and OFSE companies alike can start by asking which workflows are large enough, measurable enough, and repeatable enough to justify scaled AI deployment. A practical portfolio would prioritize use cases with four attributes: a large addressable value pool; high-quality, real-time data; clear physical feedback loops; and metrics that asset leaders already manage.

In practice, these attributes point to drilling automation, artificial lift optimization, production network optimization, waterflood and gas-lift optimization, predictive maintenance for critical equipment, real-time reliability, inventory optimization, and selected subsurface modeling applications. Gen AI productivity tools, such as document generation, coding support, and procurement assistance, can create useful efficiency, but the largest value pools are in operational performance.

For operators, the greatest opportunities can come from identifying the highest-value areas, setting standards for data and integration, making asset leaders accountable for adoption, and redesigning contracts around outcomes.

OFSE companies face a starker strategic decision over whether to be activity providers or evolve into performance partners willing to accept more risk in exchange for greater upside. The most successful OFSE players may increasingly resemble the latter, bringing proprietary workflows, automation, equipment knowledge, data infrastructure, and commercial models that translate efficiency into improved profitability for both operators and service providers.

OFSEs can also use AI to transform their own cost base. Our analysis identifies approximately $10 billion to $20 billion of internal cost efficiency potential for OFSEs across labor productivity, materials optimization, SG&A automation, and working capital improvement. This self-help to improve margins is necessary but is not sufficient to close the gap. The remainder requires commercial innovation (Exhibit 4).

Talent and operating model: Redesign workflows around decisions and accountability

For greatest value, AI can be embedded into the operating rhythm of how decisions get made: who acts, when they act, with what authority, and against what metric. In drilling, for example, this could mean moving from advisory analytics to supervised automation and eventually closedloop execution, with explicit guardrails.

In production, it could mean shifting surveillance from periodic review to full-field review with exception-based optimization, where engineers focus and act on wells and facilities with the highest expected value of intervention, and allow an AI system to manage the lower producing tail of assets. In maintenance, it could result in using AI-enabled solutions to replace calendar-based schedules with orchestrated risk-based maintenance plans tied to equipment condition and spare parts availability.

Exhibit 4

AI efficiency gains create a $24 billion value redistribution question.

Value redistribution, $ billion

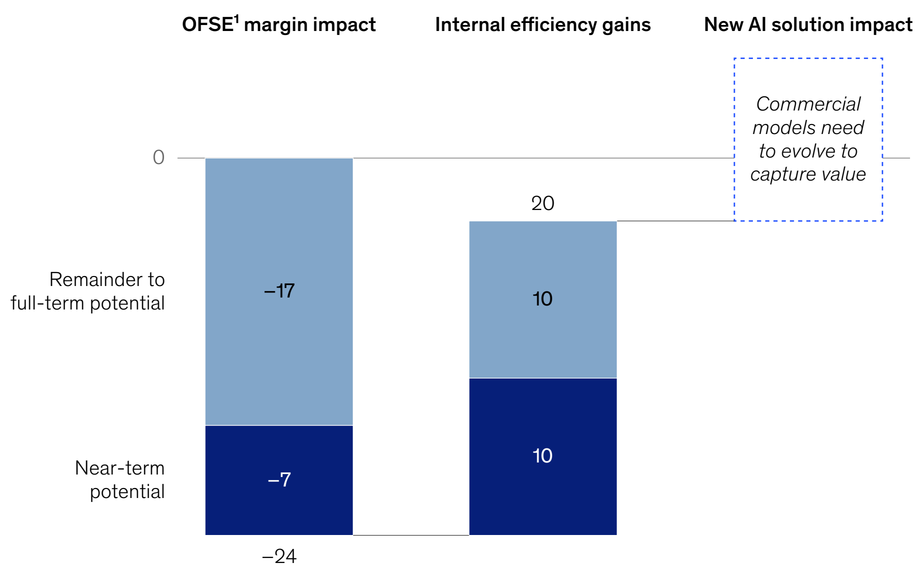

[1]Oilfield services and equipment.

McKinsey & Company

This is where agentic AI can become important. Upstream operations have many multistep workflows, such as diagnosing anomalies, checking equipment history, and verifying parts availability. Agentic systems can help orchestrate these steps, but they should be deployed in bounded workflows with clear controls and the right human-in-the-loop approvals.

Any workflow changes have direct implications for talent. The most urgent consideration is to empower teams to direct AI systems and build trust in the recommendations they create. Equally important, AI-enabled workflows should use the tacit, institutional knowledge of people, including judgment and pattern recognition, in all models and decision frameworks that must be tailored to a company's governance model and risk management protocols.

Operators and OFSE companies will need to agree on how liability is shared if AI-assisted decisions go wrong when they cross the traditional boundaries of authority and decision-making. OFSE companies cannot reasonably absorb the operational risks that may arise along the AI learning curve unless there is a proportional value-sharing mechanism in place for the upside.

Technology and data: Build the infrastructure for autonomy

The oil and gas industry has spent years building data platforms. What it now needs is infrastructure that can support operational autonomy: sensor coverage, edge connectivity, data quality monitoring, time-series architecture, model operations, cybersecurity, and governance for models that influence physical operations.

It also requires better integration between engineering and ML. Purely data-driven models can struggle in sparse or changing operating conditions, where physics-informed and deep-learning surrogate models, which combine speed with engineering constraints, are more reliable.

The goal is to have staged architecture that supports the highest-value workflows first. For example, a field with artificial lift optimization needs reliable well data, control connectivity, failure history, and a way to test and measure production uplift. A drilling automation program needs real-time rig data, surface and downhole measurements, standard event classification, and clear handoffs between automation and human supervision.

This is where the build-versus-buy-versus-partner decision matters. Where the application is commoditized, such as in predictive maintenance of rotating equipment or standard production surveillance, an OEM or service provider with fleet-scale data across thousands of installations will likely outperform any single operator's in-house effort.

On the other hand, where competitive differentiation is at stake—such as in exploration workflows, proprietary reservoir models, and trading intelligence—operators can protect their IP through internal builds or partnerships with suppliers, with clear guardrails.

Additionally, operators will need to weigh the benefit of sharing data to gain access to broader data sets across the industry against retaining data for competitive advantage. Broader training sets, accessed through joint industry projects, can mean faster AI model learning and greater accuracy, but this can be at the cost of differentiation. Operator CEOs need to make a key decision about which matters more.

Adoption: Rewire commercial models so efficiency is rewarded

Rewiring commercial models requires a paradigm shift in how the industry structures value exchange between operators and service providers. Our work points to three emerging models: AI product businesses, outcome-based pricing, and gain-sharing mechanisms. There are early examples of AI product revenue in production optimization with software-as-a-service (SaaS) solutions, as well as performance-linked models in drilling and production optimization.

We have seen model and business relationship alignment between operators and OFSE companies in the following ways:

- SaaS for repeatable workflows where the product value is measurable but operational execution remains with the operator or asset team. Baker Hughes, one of the world's largest OFSEs, offers an example: Its Leucipa™ platform optimizes production, reduces unplanned downtime, and improves efficiency.[10]

- Performance-based services where the OFSE provider influences the outcome directly, such as drilling cycle time, equipment uptime, or artificial lift performance. Consider, for example, Helmerich & Payne (H&P), a global drilling and technology company. H&P's performance contracts create win-win alignment with customers: operators get wells delivered faster and at lower cost; and H&P earns per-well bonuses for exceeding drilling speed benchmarks.[11]

- Gain-sharing models where both parties agree on baselines, measurement rules, and allocation of upside from production uplift, cost reduction, or avoided downtime.

- Over time, these models could evolve toward risk-sharing autonomy models where providers accept greater responsibility for automated decisions in exchange for higher compensation within agreed operating envelopes.

All of these models require transparent value measurement, data rights that are separated from AI model ownership, liability frameworks for autonomous decisions, capital incentives for codevelopment, and clearer build-versus-buy boundaries.

Scaling: CEOs can prioritize AI as a top agenda item for their organization
Scaling AI in a company is ultimately a CEO-led transformation. The leader's role is to move AI from a set of promising pilots to a core part of how the enterprise creates value by setting the ambition, prioritizing the selected domains and use cases, and making adoption a performance expectation. To get started, CEOs can consider the following questions:

- Where will AI create the most enterprise value, and is there an executable road map that reflects this?

- Who in the business will own AI adoption and be accountable for capturing that value?

- Are our data and technology foundations strong enough to scale solutions across operating assets?

- What will it take to make frontline adoption the default, and are our incentives and ways of working set up to get there?

---

[10] "Baker Hughes, Expand Energy to deploy Leucipa™ across US natural gas fields," Baker Hughes, January 29, 2026.
[11] "Helmerich & Payne fiscal fourth quarter 2024 earnings call," Helmerich & Payne, November 14, 2024.

The next decade: From AI experiments to capturing value

The upstream AI opportunity is significant. The value pools have been mapped and many applications are already being used, with more emerging almost weekly. The technology is improving quickly, especially in ML, physics-based models, surrogate simulation, gen AI, and agentic systems. We have identified that the largest value opportunities are concentrated in drilling, production optimization, artificial lift, reservoir management, and maintenance.

AI can make upstream faster, safer, leaner, and more productive. But the next stage will not be achieved by technology alone. It will be won by companies that can move from pilots to AI portfolios, shift from insights to bold decisions and new ways of working, build infrastructure capable of adopting operating autonomously, and develop outcome-based commercial models that share the value from AI implementation.

For CEOs ready to act, the potential path forward starts with three steps: define the AI value aspiration for your asset base and portfolio, assess organizational and technical readiness against the domains that matter most, and build then execute a sequenced road map that addresses technology, workflows, talent, and commercial models.

Giorgio Bresciani is a senior partner in McKinsey's London office, and **Bill Ambrose** and **Spandan** are partners in the Houston office, where **Priyank Singh** is a consultant.

The authors wish to thank Ignacio Fantaguzzi, Luciano Di Fiori, Nikhil Ati, and Nikhil Patel for their contributions to this article.