

From FOMU (Fear Of Messing Up) to First Win



A Practical Playbook for AI Adoption in Oil & Gas

By Kurt Kemmerly, SVP, AI Adoption & Consumption – IFS



The Stakes for Oil & Gas Operators

Oil and gas operators have spent the better part of a decade experimenting with AI, and the last two years investing in it heavily. Yet the industry's central AI question has barely moved: how to convert enthusiasm and budget into measurable operational results. Independent research suggests the gap is wide. A 2025 study from MIT's NANDA initiative found that roughly 95 percent of enterprise generative AI pilots fail to produce a measurable return, even as more than \$30 billion has flowed into enterprise AI programs ^[1].

For capital-intensive, safety-critical industries like upstream and midstream oil and gas, that failure rate is expensive in a different way than it is for a retailer or a media company: every quarter spent on an AI pilot that never reaches production is a quarter of avoidable downtime, deferred production, and unmanaged risk that a working system could have addressed ^[1].

This white paper examines why so many industrial AI initiatives stall, what the research says separates the small minority of programs that scale from the majority that don't, and lays out a practical, field-tested playbook for building early wins that compound into lasting operational value. The throughline is simple: AI adoption in oil and gas is fundamentally a sequencing, ownership, and change-management challenge, and one that can be solved with discipline.

A Familiar Question in the Operations Center, A Changed Technology

In 2017, supervised machine learning was still a novelty in oil and gas legal and commercial functions. The handful of operators willing to test it weren't applying it to production optimization or equipment monitoring - they were applying it to a different class of problem: the unstructured legal agreements that underpin upstream and midstream operations, including oil and gas leases, surface use agreements, and right of way agreements. The early customers sat in Land, Legal, and M&A rather than the operations center, and what they discovered was still durable: AI, applied to a real operational problem, could move the numbers that mattered. Those early results built the case for broader adoption across the industry.


Nearly a decade later, the technology has changed dramatically. Generative AI and autonomous agents have replaced narrow predictive models as the center of industry conversation, and capital has followed: the global market for AI in oil and gas is projected to grow from roughly \$5.1 billion in 2025 to \$18.7 billion by 2035 ^[2].

The underlying question executives ask, though, hasn't shifted nearly as much as the technology has. It is still, fundamentally: how do we move from curiosity to confidence - from a promising demo to a system operators trust and depend on?

“The question is no longer whether AI belongs in the enterprise - it's how to turn growing investment into measurable business impact. That's where many companies still struggle.”

- Kurt Kemmerly, SVP, AI Adoption & Consumption, IFS

The GenAI Divide



5% measurable value vs 95% with no measurable P&L impact.



The Adoption Gap: Why Most AI Initiatives Stall

The scale of the challenge is now well documented outside the energy sector, and the pattern maps closely onto what operators report internally. MIT's Project NANDA, drawing on more than 300 publicly disclosed enterprise AI deployments, 52 organizational case studies, and 153 leadership interviews, found that only about 5 percent of generative AI pilots were extracting measurable financial value^[1]. The remaining 95 percent showed no measurable effect on profit and loss, despite genuine executive commitment and real budget behind them^[1].

The researchers concluded the divide has less to do with model quality than with integration. General-purpose generative tools perform well for individual users precisely because they are flexible, but that same flexibility makes them brittle inside enterprise workflows, where they cannot retain institutional context, learn from operator feedback, or integrate with the systems (ERP, historian, SCADA, maintenance management) that actually run the business^[1,3]. Tools that stay outside those workflows tend to look polished in a demonstration and then fail to survive contact with day-to-day operations^[1,3].

A second, related finding is a misalignment of spend. Roughly half of enterprise generative AI budgets, in MIT's sample, went to visible, high-profile functions such as sales and marketing, while some of the clearest and most repeatable returns showed up in less glamorous, workflow-embedded automation - the oil and gas equivalent of production surveillance, payables processing, and maintenance scheduling rather than customer-facing pilots^[1].

For oil and gas operators, the practical takeaway is that most organizations are testing AI in the wrong place, in the wrong sequence, and without the operational ownership needed to carry a pilot into daily use.

The Cost of Standing Still

While companies debate pilot design, the operational cost of the status quo continues to accrue. Unplanned downtime is one of the largest and most measurable of these costs. Estimates vary by methodology and region, but they consistently point to material sums: industry surveys put the average cost of unplanned downtime in oil and gas operations at roughly \$220,000 to \$500,000 per hour depending on facility type, with offshore assets often exceeding \$1 million per hour of lost production^[4,5].

At the site level, some analyses put the annual cost of unplanned downtime at well over \$100 million per

facility once direct production losses, emergency repair costs, and secondary effects are included, with one industry study noting downtime costs at oil and gas sites rising more than 75 percent over the preceding several years^[6,7].

Predictive, AI-enabled maintenance is one of the few levers with a well-established track record against this problem. An analysis of industrial AI deployments by McKinsey, the global management consultancy, found that predictive analytics can lower maintenance costs by as much as 25 percent while improving asset availability. Their analysis also found that integrated AI systems deployed at scale, rather than as isolated pilots, can reduce operating expenditures by up to 20 percent and lift production efficiency by 5 to 8 percent across upstream and downstream operations^[8]. Boston Consulting Group has documented similar magnitude effects, with AI-driven operational programs shortening key cycle times from months to weeks and reducing operating costs by an estimated 15 to 20 percent in capital-intensive environments^[8].

These figures matter for a specific reason: they establish that the value AI can create in oil and gas operations is not speculative. The technology has a demonstrated ceiling well above what most organizations are currently capturing. The gap between that ceiling and today's results reflects an adoption problem more than a capability problem.

What the Successful Minority Do Differently

If the technology itself is not the primary constraint, what distinguishes the organizations that do capture measurable value? The research points to three consistent factors: narrow initial scope, deep workflow integration, and, perhaps most decisively, active executive sponsorship paired with clear business ownership.

Narrow Scope Beats Broad Ambition

MIT's research found that the AI programs generating real value tended to start from a single, well-defined operational pain point rather than an enterprise-wide transformation agenda. Analysts studying the same pattern in industrial settings have reported that companies deploying AI at scale around a specific, high-value use case saw measurable operational improvements of 15 to 20 percent within 12 to 18 months, while broader, less targeted initiatives were far more likely to stall before reaching production^[9].

Sponsorship Is the Strongest Predictor of Success

Outside the AI literature specifically, decades of change-management research point to the same

conclusion from a different angle. Benchmarking studies by Prosci, a change management consultancy, are the most extensive longitudinal research on organizational change, running continuously since 1998. Those studies identified active, visible executive sponsorship as the single greatest contributor to successful change in every study conducted ^[10].

Organizations with strong, visible sponsorship report success rates around 73 percent, compared with roughly 29 percent for initiatives lacking it, and projects with highly effective sponsors are close to three times more likely to meet their objectives than those with ineffective ones ^[10].

This is directly relevant to industrial AI, which is not simply a software rollout. It changes how operators make decisions, what they trust, and who is accountable when a model's recommendation conflicts with experience. Prosci's own recent research into AI-specific adoption reinforces that AI programs require a change strategy distinct from a standard technology deployment, and that coalition-building by sponsors is one of the strongest predictors of sustained employee engagement with new AI tools ^[11].

Ownership Has to Sit With the Business, Not Just IT

The organizations MIT studied that did break through the “GenAI Divide” shared a structural trait: accountability for outcomes sat with a business owner close to the operational problem, not solely with a technology or data science function. That owner was responsible for the metric the AI system was meant to move - hours of downtime avoided, barrels protected, cases resolved - rather than for the technology's deployment alone ^[1].

A Practical Playbook for First Wins

Bringing these threads together - the operational cost data, the MIT research on what fails, and decades of change-management evidence on what succeeds - points toward a repeatable sequence for oil and gas operators. None of the five principles below is exotic. Their power comes from being applied together, and applied early.

1. **Start with a business champion, not a technology mandate.** Identify an operational leader who owns a specific pain point and who is motivated to prove a result, not simply to sponsor a project. Prosci's research shows a project's odds of meeting its objectives climb from roughly one in four to better than four in five when sponsorship is active and visible rather than nominal ^[10].
2. **Pick one high-value use case and prove it in weeks, not quarters.** Resist the instinct to build an AI “program” before there is a single working result. A narrow scope - one production system, one failure mode, one workflow - is easier to instrument, easier to validate against ground truth, and easier for skeptical operators to trust once it works.
3. **Measure outcomes operators already track, not model metrics.** Hours of downtime avoided, production protected, incidents prevented, and cost per barrel are the language of the business. Model accuracy and precision-recall curves speak a different language entirely. Programs that report in operational terms build credibility faster than programs that report in data-science terms.
4. **Build visible momentum inside the first 90 days.** The first quarter sets the tone for whether an initiative is seen as a pilot that will quietly disappear or a capability the organization is building toward. An early, well-communicated result, even a modest one, does more to overcome organizational skepticism than a technically superior result delivered a year later.
5. **Pair executive sponsorship with accountable business ownership.** Sponsorship opens doors and protects resources; ownership makes sure the work gets done and the metric moves. The research above shows these are complementary: programs need both a leader who removes obstacles and an owner who is answerable for the number.

Applied consistently, this sequence replaces uncertainty with evidence. Each proven use case becomes the reference point and the internal case study that justifies the next one, turning what began as a single pilot into an operating discipline.

From Pilot to Practice: Sustaining Momentum

A first win is necessary but not sufficient. The MIT and Prosci research both point to the same failure mode after an initial success: organizations that treat the first proof point as the finish line rather than the template. Scaling requires the same discipline applied a second and third time - new business champion, new narrow use case, same measurement rigor, same sponsorship model - rather than an attempt to generalize the first system into a broader platform before it has been proven elsewhere.

It also requires resisting the pull of prestige projects. As MIT's researchers observed, enterprise AI budgets have skewed toward visible, customer-facing applications even though some of the most reliable returns have come from less visible, workflow-embedded automation ^[1]. For an oil and gas operator, that suggests looking first at production surveillance, maintenance scheduling, and back-office processes with clear, auditable outcomes: the operational equivalent of the “boring” use cases MIT found were quietly outperforming the more visible ones ^[1].

Replacing FOMU With Evidence in the Field

Fear of messing up - the FOMU in this paper's title - is a rational response to a genuinely difficult track record. A 95 percent failure rate for enterprise AI pilots is the documented experience of most organizations that have tried ^[1]. But the same research that quantifies the failure rate also identifies what the successful minority did differently, and none of it depends on being first, having the largest budget, or building the most sophisticated model.

It depends on choosing a real operational problem, assigning a business owner who is accountable for a result operators care about, pairing that owner with active and visible executive sponsorship, and proving value quickly enough that the organization can see the evidence rather than take it on faith. For an industry built on managing risk under uncertainty, that is a familiar discipline applied to a new technology - and, based on the available evidence, it is the discipline that reliably turns AI investment into operational results.

Citations & Sources

1. Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). The GenAI Divide: State of AI in Business 2025. MIT NANDA. <https://nanda.media.mit.edu>
2. Market.us. (2026). Global AI in Oil and Gas Market Size, Share, CAGR Forecast 2026-2035. <https://market.us/report/global-ai-in-oil-and-gas-market/>
3. Snyder, J. (2025, August 26). MIT Finds 95% of GenAI Pilots Fail Because Companies Avoid Friction. Forbes. <https://www.forbes.com/sites/jasonsnyder/2025/08/26/mit-finds-95-of-genai-pilots-fail-because-companies-avoid-friction/>
4. MarketScale. (2021). What's the True Cost of Machine Downtime? <https://marketscale.com/industries/engineering-and-construction/whats-the-true-cost-of-machine-downtime/>
5. Berisintl. (2025). The Real Cost of Equipment Downtime for Oilfield Operations. <https://www.berisintl.com/the-real-cost-of-equipment-downtime-for-oilfield-operations>
6. Innovapptive. (2024). Solving Unplanned Downtime Challenges in the Oil & Gas Industry. <https://www.innovapptive.com/blog/challenges-of-oil-gas-equipment-maintenance>
7. iFactory. (2025). The Hidden Cost of Unplanned Downtime in Manufacturing. <https://ifactoryapp.com/blog/hidden-cost-unplanned-downtime-manufacturing>
8. Usetech. (2026). From the Well to the Neural Network: The Future of Oil & Gas with AI in 2026 (citing McKinsey & Company and Boston Consulting Group analyses). <https://usetech.com/blog/from-the-well-to-the-neural-network-the-future-of-oil-gas-with-ai-in-2026/>
9. Apidots. (2026). Top 7 AI Uses in Oil & Gas Industry (2026). <https://apidots.com/blog/ai-uses-oil-gas-industry/>
10. Prosci. (2025-2026). Best Practices in Change Management: Top Contributors to Success (benchmarking research series, 1998-present). <https://www.prosci.com/blog/change-management-best-practices>
11. Prosci. (2026). Keys to Unlocking AI Adoption (State of AI adoption research). <https://www.prosci.com/blog/ai-adoption>