WHITE PAPER

AI Capability & Workforce Readiness in Energy and Industrial Enterprises

From Pilots to AI-native Operations

Asaad Abdollahzadeh · Founder & Managing Director, Zanor AI · 22 April 2026

Executive Summary

Energy enterprises have spent the last three to five years investing heavily in artificial intelligence. The evidence is everywhere: proof-of-concept notebooks, predictive maintenance pilots, drilling optimisers, reservoir characterisation models, digital twin experiments, generative AI copilots for engineers, computer vision for inspection, LLM assistants for planning and operations. Yet when leaders audit the operational impact of that investment, the pattern is almost always the same. A handful of use cases deliver measurable value. Most sit in a state of perpetual pilot. Few are owned, operated, and evolved by the business itself.

The prevailing explanation blames the technology, the data, or the platform. It is the wrong explanation. The consistent failure mode is not technical. It is organisational, and it is compounded by a second fact that is easy to overlook: AI and data systems are not traditional software. They behave differently, fail differently, cost differently, and require a different operating discipline. Energy enterprises that deliver AI using their traditional IT playbook reliably produce orphaned pilots.

This paper sets out Zanor AI's perspective on what it takes for energy enterprises to move beyond fragmented pilots and become AI-native organisations, with humans firmly in control. It clarifies how AI systems differ from traditional software and why that difference reshapes capability requirements. It introduces a practical capability framework covering three dimensions (technical, operational, leadership) and four maturity levels (awareness, adoption, integration, AI-native). And it argues that the critical missing layer in almost every AI programme today is structured, practitioner-led workforce readiness with domain experts at the core of delivery from day one.

Four recommendations anchor the paper

1. Treat AI as an operating model transformation, not a technology one. The receiving organisation matters more than the algorithm.
2. Recognise that AI systems are fundamentally different from traditional software. Their lifecycle, economics, and failure modes require a different operating discipline.

3. Build capability through real energy use cases, not classrooms. Make programmes practitioner-led and embed learning in delivery.
4. Put domain experts in the core team from day one. SMEs are co-owners of AI systems, not stakeholders consulted occasionally.

1. The Current State of AI in Energy

Across upstream oil and gas, LNG, refining and downstream, power generation, transmission and distribution, and renewables, AI adoption has moved past curiosity. Investment is significant. Ambition is higher still. But operational penetration remains thin.

Three patterns recur across the sector:

Fragmented pilots with limited reuse

Multiple assets, basins, or business units independently commission AI work with different vendors, on different platforms, using different data conventions. A rotating equipment monitor built for one compressor train rarely moves to the next. A drilling optimiser developed in one basin stays in that basin. A grid forecasting model built by one regional utility operator cannot be replicated in a neighbouring one. Each pilot is a small island.

Limited scale and sustainability

Even successful pilots struggle to move into production. Models decay as operating conditions change, dashboards go stale, and when the original delivery team leaves, the system quietly falls out of use. The organisation is left with screenshots in a steering committee deck but little operational value in the control room, planning meeting, or shift handover.

Over-reliance on vendors and external partners

Because internal teams were not part of the build, they cannot operate the system. Every change, every retrain, every extension requires a procurement cycle. The asset exists on paper; the capability to evolve it does not.

The result is a widening gap between AI investment and AI impact. The problem is not that the industry lacks good use cases or competent vendors. The problem is that the receiving organisation is not structurally ready to absorb what is being delivered, and has not recognised that AI systems require an operating discipline that is different from the one it has used for decades of traditional software delivery.

2. Why AI Systems Are Not Traditional Software

Energy enterprises have decades of muscle memory running traditional software projects: gather requirements, design, build, test, deploy, run. Applying that same playbook to AI systems is a common and costly mistake. AI and data systems differ in kind, not degree, and the differences shape almost every capability decision that follows.

Seven differences matter most:

Dimension	Traditional Software	AI and Data Systems
Behaviour	Deterministic. Same input, same output.	Probabilistic. Output quality depends on data, context, and operating conditions.
Lifecycle	Ages gracefully. Runs until changed.	Decays silently. Requires monitoring, retraining, and retirement.
Primary asset	Code is the IP. Data is input.	Data is the IP. The model is a compiled expression of it.
Testing	Unit and integration tests against specifications.	Continuous evaluation against moving ground truth and operational outcomes.
Economics	Fixed licence and run cost.	Variable inference cost, retrain cost, and LLM token economics.
Starting point	Requirements specification.	An operational decision worth improving.
Delivery team	Software engineers can build to a spec.	Multi-disciplinary. Domain expert is a core builder, not a reviewer.

Why this matters for energy

These differences are not abstract. They are the reason AI projects in energy fail in ways that traditional IT projects do not.

A compressor anomaly detector built to a traditional software spec will treat every model output as correct, regardless of whether the unit has entered an operating mode the model never saw in training. A drilling parameter recommender built without continuous evaluation will quietly drift as lithology and rig conditions change. An LLM assistant deployed to planners without token cost visibility will generate a six-figure surprise on the first month's invoice. None of these are technology failures. They are the predictable consequences of running AI systems with a traditional software mindset.

A mature AI operating discipline recognises that AI systems are living systems. They must be designed, monitored, retrained, and retired deliberately. The capability requirements, operating model, and governance all follow from that single recognition.

3. Why AI Initiatives Fail to Scale

When Zanor AI looks across AI portfolios in energy, five failure modes account for the majority of stalled programmes. None of them are solved by better algorithms.

Lack of internal ownership

No one inside the organisation has a clear mandate to operate the model after go-live. Accountability ends where the vendor statement of work ends, and the model becomes orphaned.

Weak operating model

The roles, decision rights, and workflows required to run AI in production (model owner, data steward, value owner, MLOps engineer, governance committee) are missing or informal. The programme runs on heroics, not on structure.

Skills mismatch across business and technical teams

Data scientists do not understand the asset, the process, or the decision the model is supposed to improve. Engineers and operators cannot interrogate a model's behaviour. Leaders cannot tell a good AI investment from a bad one. The language barrier prevents the conversation from happening at the right altitude.

Domain expertise brought in too late

The subject matter expert is consulted occasionally rather than embedded from day one. The team builds what is easy, not what is needed. In energy, where process physics, operating envelopes, and safety constraints shape what is acceptable, late SME involvement is the single most consistent predictor of AI project failure.

Poor integration into operational workflows

The model produces an output; the decision is still made the way it was made before. Without deliberate changes to control room practice, maintenance scheduling, planning meetings, or shift handover, AI outputs are curiosities, not decisions.

Absence of governance and lifecycle management

There is no systematic view of which models are in production, how they are performing, when they last retrained, or whether they remain fit for purpose. Risk accumulates silently until something breaks.

CASE STUDY · THE POC GRAVEYARD

An international energy operator accumulated more than one hundred proof-of-concepts (POCs) across upstream, midstream, and downstream over three to four years. A centralised internal review found that only a small fraction had been converted into full-scale, in-production use cases.

The common pattern behind the stranded POCs was not model quality. Most POCs had been scoped and commissioned by central engineering, digital, or technology functions, or by vendors, not by the frontline operational teams that would have to use them. No business owner from the operational or end-user team had been designated as a co-owner of the model from day one. No serious data readiness assessment had been performed. No operational role had been redesigned to use the output. When POC funding and timelines ended, no one was accountable for taking the work to the next stage.

The fix was not more POCs. It was a practitioner-led workforce readiness programme that embedded domain experts in the core delivery team and redefined the subject matter expert role to own and evolve the use case alongside the AI model built for it.

4. The Missing Layer: AI Capability and Workforce Readiness

AI capability is not a headcount metric. It is not solved by hiring ten more data scientists. A capable energy organisation is one where the right people, at the right levels, have the right skills and the right accountability to use AI to improve operational and business decisions, and to sustain that improvement over time.

Zanor AI defines AI capability across three interlocking dimensions:

Technical capability

The ability to build, deploy, and maintain AI systems in an energy context. This covers data engineering against industrial historians and sensor networks, machine learning and LLM engineering, platform and MLOps/LLMOps, security, integration with control and operations systems, and the engineering discipline to manage AI as a living system rather than a one-off build.

Operational AI literacy

The ability of reservoir, drilling, production, process, rotating equipment, and grid engineers, along with operators, planners, and traders, to work with AI outputs as part of their daily decisions. Literacy is not about writing code. It is about knowing when to trust a model, when to challenge it, what a drift warning means, and how to convert a model output into a better operational decision.

Leadership understanding and sponsorship

The ability of executives and functional leaders to set AI strategy, allocate investment, govern risk, and hold the organisation accountable for value. Leadership capability is what turns AI from a collection of projects into an enterprise capability.

All three dimensions must move together. An organisation strong on technical capability and weak on operational literacy will build models nobody uses. An organisation strong on leadership and weak on technical depth will buy platforms it cannot run. An organisation that excludes its domain experts from the core team will build AI that is technically competent and operationally wrong.

5. A Practical AI Capability Framework

Zanor AI uses a three-part framework to diagnose an energy enterprise's AI capability position and to design the path forward. The framework is deliberately practical. It is designed to be used in a workshop with the executive team, not to be admired in a report.

5.1 Capability Dimensions

Three dimensions, each with concrete sub-capabilities. An energy enterprise needs progress in all three to reach AI-native operations.

Three dimensions — progress is required in all three to reach AI-native operations.

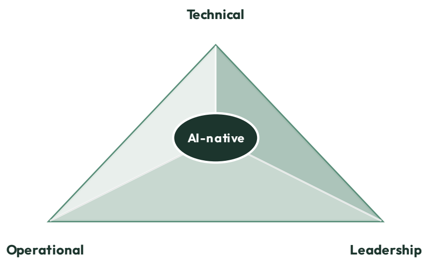

Technical
- Industrial data foundations
- ML and LLM engineering
- Platform, MLOps and LLMOps
- Security, OT integration
- Living-system engineering

Operational
- Use-case shaping with SMEs
- Decision integration
- Model consumption literacy
- Continuous feedback loops
- Frontline change capability

Leadership
- AI strategy for energy
- Portfolio governance
- Value realisation
- Risk, safety, and ethics
- Change leadership

5.2 Workforce Segmentation

Capability means different things to different people. A credible programme addresses four distinct workforce segments. Generic AI training programmes collapse these segments into a single curriculum. That is why they produce limited impact.

Segment	What they need to do	What capability looks like
Executives	Set direction, govern risk, allocate capital, sponsor change	Fluency in AI economics, governance, strategic use cases, and value tracking
Domain experts (reservoir, drilling, production, process, rotating equipment, grid)	Co-own use cases from day one, challenge data and models, embed outputs into daily operational decisions	Operational AI literacy, confidence to critique and improve models, clear decision rights as model co-owners
Data and AI practitioners	Build, deploy, monitor, and evolve AI systems in production alongside SMEs	Full stack: data, ML and LLM engineering, MLOps/LLMOps, energy domain fluency
IT, OT, and platform teams	Provide and run the AI platform, security, IT/OT integration, and scale	Platform engineering, MLOps, OT-aware security, AI-aware architecture and cost control

5.3 Capability Maturity Levels

Four maturity levels describe where an enterprise sits and where it needs to go. Most energy enterprises today sit between Awareness and Adoption. The step-change in value comes between Adoption and Integration, when AI stops being a set of projects and becomes part of the operating fabric, and when domain experts move from reviewers to co-owners of AI systems.

Capability maturity levels — most enterprises sit between Awareness and Adoption today.

4 **AI-native** *"AI is how we operate."*	AI integral to decision-making. Organisation evolves systems independently.
3 **Integration** *"AI is part of how we operate."*	AI embedded in operational workflows. SMEs in core teams. Lifecycle management in place.
2 **Adoption** *"We have AI in a few places."*	Selected use cases in production. Technical team forming. Governance emerging.
1 **Awareness** *"We are exploring AI."*	Pilots exist, mostly vendor-led. Limited internal understanding. No operating model.

6. Designing Targeted Capability Programmes

Most AI training programmes in the industry are generic, classroom-heavy, and divorced from the organisation's actual use cases. They generate certificates, not capability. Effective capability programmes for energy enterprises share four attributes.

Role-based, not generic

Content, depth, and delivery format are tailored to the workforce segment. An executive does not need to understand backpropagation; an MLOps engineer does. A plant operator needs to know how to interrogate a model's output during a shift; a data scientist needs to know how to build and monitor that model. A reservoir engineer needs to know how to challenge a geosteering model; a drilling supervisor needs to know when to override it.

Use-case driven

Learning is anchored in the enterprise's actual operational problems, not hypothetical Kaggle datasets. A capability programme for a refinery uses refinery data, refinery decisions, and refinery KPIs. A programme for a grid operator uses load, weather, and asset telemetry from its own network. This forces relevance and produces reusable operational assets in parallel with building skills.

Practitioner-led, not trainer-led

Capability is transferred by people who have built, deployed, and operated AI systems in energy. Classroom trainers cannot teach what they have not done. Academics can teach theory but struggle with the messy reality of production AI in an operating environment: data drift as wells age, sensor failures in harsh conditions, model overrides in safety-critical contexts, regulatory scrutiny, legacy control system integration. A credible programme is led by practitioners who have lived these problems, often for years, and can translate the craft. The line between delivery and capability building disappears. The same people who are building the model are building the team.

CASE STUDY · WHEN CAPABILITY LOSES MOMENTUM

A national energy operator launched a Citizen Analytics programme to build AI, analytics, and data capability across enterprise, engineering, and operations. Early traction was strong, with engaged teams and several initial use cases delivered. Within one to two years, momentum declined — not due to lack of capability.

It was drift from core principles. Use cases became less tied to real operational problems. Delivery shifted from practitioner-led to trainer-led. Early use cases were not taken through to tangible operational value, even when small. The programme built awareness, but not ownership.

The reset was clear: return to use-case-driven learning, embed practitioners in delivery, and ensure every use case delivers visible value. Capability only sticks when it is built through real work, not training alone.

Domain experts in the core team from day one

The most consistent predictor of AI success in energy is whether the domain subject matter expert is in the core team from the first day, not a stakeholder consulted occasionally. Reservoir, drilling, production, process, rotating equipment, and grid engineers belong in the room for every significant design choice: which features matter, what operating modes must be covered, where the physics constrains the model, what operator action the output will drive, and how to recognise when the model is about to mislead a decision. A credible programme treats the SME as a co-owner of the system, not a reviewer. In practice, this means the SME attends daily stand-ups, contributes to model evaluation, signs off on releases, and is recognised and resourced accordingly.

The three elements of a well-designed programme

A well-designed programme combines three elements: hands-on use-case development with internal teams progressively leading more of the work; co-delivery and structured shadowing across technical, operational, and leadership layers; and formal knowledge transfer artefacts including playbooks, runbooks, decision guides, governance templates, and model cards.

The test of a capability programme

The test of a capability programme is not how much people learned. It is whether, six months after the external team leaves, the organisation is still building, operating, and evolving AI systems independently.

PEER PERSPECTIVE · THE INSTITUTIONAL IMMUNE RESPONSE

In a recent LinkedIn post, Prakash Kumar Karunakaran described an LLM completing in two minutes a calculation that had previously taken him three days, and how institutional resistance blocks that kind of capability from scaling across an engineering workforce. His observation: only a small fraction of organisations generate real AI value, and the ones that do succeed replace generic training with domain-specific application.

Responding to that post, Zanor AI's founder added: the issue was never the technology. It is how organisations react to it. The institutional immune response slows adoption. Domain-grounded, practitioner-led application, not generic training, is what drives meaningful AI integration and competitive advantage.

This paper is written with that experience in mind.

7. Operating Model and Governance as Enablers

Capability without operating model evaporates. Trained people, without clear roles, decision rights, and processes, revert to old habits. Four operating model elements are non-negotiable for scaled AI in an energy enterprise:

AI governance

A defined body with clear scope: use-case prioritisation, risk review, model approval, value tracking, and ethics. The governance body sits close enough to the business to make real decisions and close enough to delivery to act on them. In energy, it must connect to existing HSE, engineering assurance, and operational risk structures.

Model lifecycle management

Explicit processes for how models are proposed, built, validated, deployed, monitored, retrained, and retired. This is the industrial discipline that keeps AI reliable at scale. It is the practical expression of the recognition that AI systems are living systems, not static software.

Cross-functional collaboration

Standing interfaces between the AI team, the operational business, IT and OT, security, and risk. Without these, every project rediscovers the same friction and the same delays, and the same domain experts are re-recruited into each new pilot from scratch.

Clear ownership and accountability

Each production model has a named business value owner, a named technical owner, and a named domain expert co-owner. All three are measured on the model's operational contribution. All three are resourced to act.

Governance is often framed as a brake on innovation. In well-run programmes it is the opposite. It is the reason innovation reaches production and stays there.

INDUSTRY OBSERVATION · FROM VENDOR-LED PLATFORM TO INTERNAL OWNERSHIP

In the late 2010s, Equinor made a deliberate shift from vendor-led advanced analytics to an in-house data and AI capability. The company stood up Omnia, its enterprise data platform on Microsoft Azure, and built a central data science organisation of several hundred practitioners embedded alongside subsurface, drilling, and production engineering. Publicly, Equinor has framed the direction as owning the data and the models that the business depends on, while still engaging vendors for acceleration on specific use cases.

The reported rationale was speed and defensibility. When every model change had to travel through a vendor statement of work, the pace of model evolution could not keep up with the operational cycle. Internal ownership of the platform, the data, and the model portfolio let Equinor iterate on reservoir, drilling, and production models in step with its own engineering decisions, and explain those models to regulators without a third party in the room.

Sources: Equinor press releases on the Omnia platform launch, 2018 onwards; Microsoft Customer Stories (Equinor); public talks by Equinor data leadership at Microsoft Ignite and Databricks Data + AI Summit.

8. From Capability to Independence

The strategic destination is not self-sufficiency in the sense of doing everything alone. It is the confidence that the energy enterprise owns the AI capability it relies on, and can choose when to use partners and when to build internally.

Independence has three hallmarks:

- Owning AI systems internally. Models, data pipelines, platforms, and the knowledge required to run them sit with the organisation, not with a vendor.
- Reducing dependency on external vendors for run-the-business work. Partners are used for acceleration, scaling, and new capability, not for keeping the lights on.
- Continuously evolving models and systems. The organisation has the in-house capability, the operating model, and the investment discipline to evolve AI systems as reservoirs deplete, processes are reconfigured, fleets age, and the energy mix shifts.

This does not mean eliminating partners. It means using them intentionally. Zanor AI's delivery model is explicitly designed to make itself progressively unnecessary for run-the-business activity, while remaining a partner for the next wave of capability.

INDUSTRY OBSERVATION · OPERATIONAL LITERACY AT THE FRONTLINE

Shell's multi-year AI upskilling programme, delivered with Udacity, trained more than two thousand employees across technical, operational, and leadership functions in applied AI and data science. The programme was not built as generic training. It was anchored in Shell's own use cases across predictive maintenance, subsurface imaging, and energy trading, with internal practitioners co-teaching alongside external instructors.

Shell has publicly attributed much of the business value from its AI investments to the graduates of this programme, who brought operational context, process knowledge, and physical-asset understanding that external data scientists could not replicate. The industry pattern is consistent: the deepest value comes not from importing data scientists, but from equipping the people who already understand the asset to also work with AI.

Sources: Udacity customer story on Shell; Shell.ai public communications; Daniel Jeavons (Shell Digital Technologies) and Shell team public conference talks at Databricks Data + AI Summit and Spark Summit.

9. Expected Outcomes and Business Impact

When capability and workforce readiness are treated as the delivery backbone, the outcomes compound across three horizons.

Internal teams become capable of

- Owning AI systems end to end, from data to decision.
- Operating, monitoring, and maintaining models in production as living systems.

- Scaling successful use cases across assets, basins, fleets, and grid regions.
- Challenging, improving, and retiring models as operating conditions change.

The business outcomes that follow

- Improved operational efficiency: fewer unplanned shutdowns, better asset utilisation, optimised energy and chemical consumption, higher recovery factors.
- Faster decision-making: control room, planning, and maintenance decisions move from weekly cycles to near real-time.
- Increased reliability and performance: models are trusted because they are governed, maintained, and co-owned by the engineers who use them.
- Sustainable AI adoption: value continues to grow after the initial programme ends, rather than peaking at go-live.

The important shift is from individual pilot value to portfolio value. AI-native energy enterprises do not rely on any single use case. They rely on the capability to conceive, deliver, and operate a continuous stream of use cases.

10. Conclusion: The Shift to AI-native Energy Enterprises

> The difference is not measured in models deployed. It is measured in decisions improved.

The defining constraint on AI value in energy is not the technology. It is the organisation that receives it.

AI is not a technology transformation. It is an operating model transformation, delivered through a workforce transformation, built on the recognition that AI systems are fundamentally different from traditional software. Capability and workforce readiness are the foundation. Operating model and governance are the structure. Technology is the tool. In that order.

Enterprises that accept this ordering will move from AI-assisted, where a few teams experiment with AI on the side, to AI-native, where AI is embedded in the decisions that run the business, and humans remain firmly in control.

Zanor AI works with energy enterprises to design and deliver this shift. The starting point is usually a capability diagnostic: where is the organisation today, across technical, operational, and leadership dimensions, for each workforce segment, and how deeply are domain experts embedded in AI delivery? From there, a targeted, practitioner-led capability programme, anchored in real energy use cases and embedded in delivery, builds the foundation for AI-native operations.

The question for leaders is not whether to invest further in AI. It is whether the organisation receiving that investment is ready to absorb it, operate it, and evolve it. That readiness is what separates the energy enterprises that will operate AI-natively in the next cycle from those that will still be counting pilots.

ABOUT ZANOR AI

Zanor AI is a specialist consultancy focused on AI strategy, delivery, and capability for energy enterprises. The firm partners with CIOs, CTOs, Chief Digital Officers, and operational leaders to design AI operating models, deliver high-value use cases, and build durable internal capability. Zanor AI's engagement model is practitioner-led and built around progressive transfer of ownership, with domain experts embedded in the core delivery team from day one. The goal is always an internal team that can independently sustain and scale what has been built.

More practitioner papers are in our **Insights** library.