ENERGY

A Real-World Game Plan for AI in Renewable Energy

By Christopher Tuot, Anita Oh, Frank Klose, Andrej Levin, and Vinoj Pillai

ARTICLE JUNE 18, 2026 8 MIN READ

Almost all large companies today are eyeing the benefits of AI-driven transformations and increasing investment in the technology. But energy and utility players are accelerating spending faster than most. According to BCG's latest AI Radar survey, these companies plan to triple their AI spend this year compared with 2025—a bigger ramp-up than for all other sectors except insurers.

Despite the growing enthusiasm, integrated energy companies and pure players are struggling to create tangible value from AI in their renewables businesses. While companies often run AI pilots, they do not scale them effectively. Initiatives tend to remain confined to digital teams or high-level roadmaps, lack a measurable impact on the business, or fail to deal with issues on the front line. As a result, few players are seeing any significant operational or financial enhancements from AI.

The Value Creation Potential from Scaling AI in Renewable Energy

Using AI-enabled solutions in renewables creates substantial value across the business, from operations and site selection to support functions like procurement. For example, we estimate that worker productivity can be increased by 15% to 25%, while energy yield can be boosted by one to three percentage points through better asset availability. (See Exhibit 1.) What's more, managers don't have to implement and oversee hundreds of different use cases. Our experience has shown that by building 10 to 15 use cases, companies typically achieve 60% to 70% of the overall value potential. (See Exhibit 2.)

EXHIBIT 1

Deploying AI in Renewables Can Boost Productivity and Increase Energy Yield

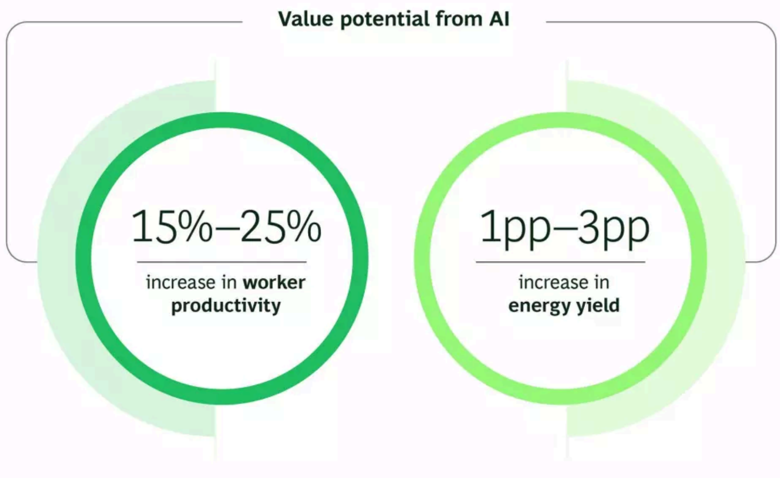

Source: BCG analysis.

EXHIBIT 2

Top 10 to 15 AI Use Cases Already Capture 60% to 70% of Value Potential

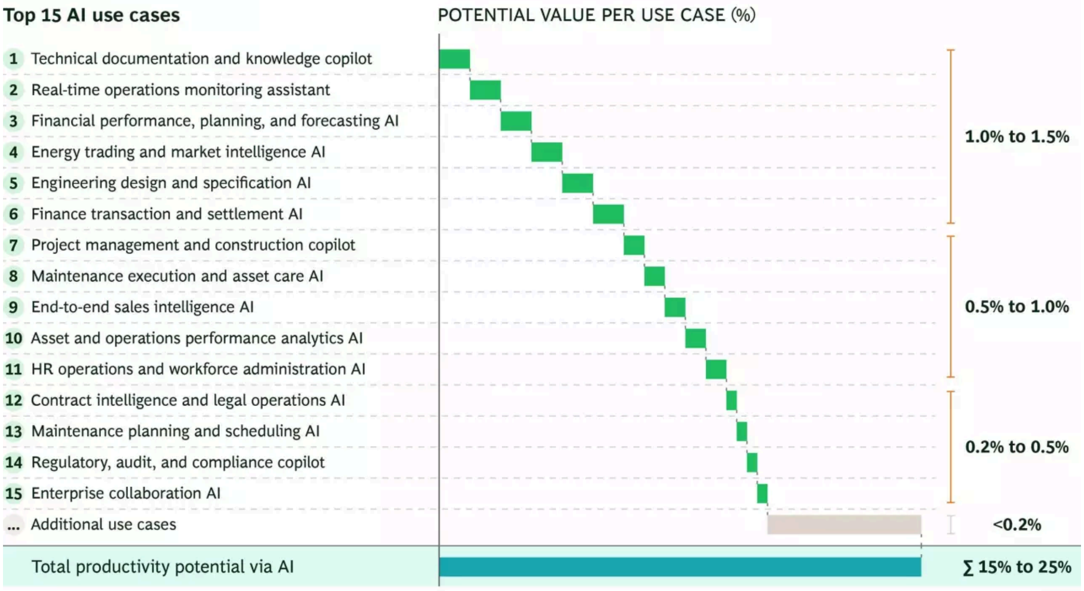

Source: BCG analysis.

However, renewable energy companies must act with care and determination when implementing AI. Rising costs and policy headwinds have contributed to margin compression in the sector in recent years. Consequently, AI investments need to deliver a quick payback. Companies also have to contend with multiple external uncertainties relating to wholesale energy prices, connection timelines, and curtailment risk. So, using AI to manage internal costs is paramount.

Subscribing to the latest, cutting-edge GenAI models alone won't fix the problem of how to get the most out of AI. Companies need to go further and adopt a structured approach to AI implementation. Based on our work with clients, we've identified six lessons to ensure that scaling AI in the energy industry delivers a meaningful impact for renewables players. These include maintaining a razor-sharp focus on value, solving operational constraints, and embedding solutions into actual workflows.

Lesson 1: Focus on value creation from the outset, and establish KPIs quickly.

Value in renewables can be viewed through the lens of performance metrics, particularly energy yield and asset availability, as well as measures across the value chain, such as the cost of operations per megawatt-hour or capex per kilowatt-hour with planned projects.

Rather than starting their AI journey with a grand ambition or end-to-end transformation, companies should aim at achieving a material improvement in a significant metric within 6 to 12 months. Link each AI initiative to a single key performance indicator (KPI) or a handful of KPIs that matter for a wind turbine or solar farm—and show how the use case moves the dial at a plant or P&L level for that asset before scaling. If you cannot clearly state which KPIs you will be measuring and explain the mechanism for financial impact in week one, you are not running an AI initiative, you are conducting research. (See Exhibit 3.)

EXHIBIT 3

Six Lessons for Successful AI Adoption in Renewable Energy

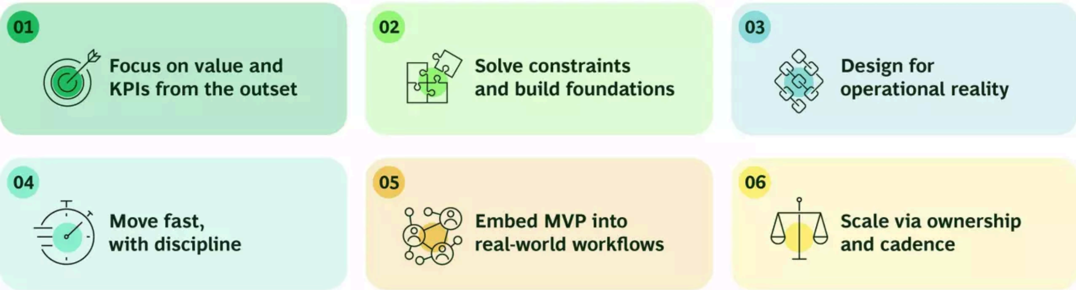

Source: BCG analysis.

Lesson 2: Solve constraints while building the foundations for other applications.

Implementing AI in asset-intensive industries such as renewable energy comes with unique challenges. Value is rarely created through beautifully redesigned processes; it is achieved by solving operational constraints so that the business can function more effectively. To maximize AI's potential, companies should prioritize creating a small slice of value before considering the implications for other parts of the business. AI applications should both solve a specific bottleneck—such as ordering parts for maintenance activities or optimizing works delivery schedules—and have a significant impact. Concentrate on fixing the constraint that is holding back your KPI to create measurable value.

The aim of AI is not to develop a set of disconnected projects, but to foster a virtuous cycle where one initiative helps build the foundations for the next. Consequently, beyond fixing a specific bottleneck, companies also need to think in terms of components that can be reused elsewhere in the business and design the initiative with them in mind. These components include ownership responsibilities, governance approaches, integration mechanisms, and data structures. By starting with a targeted, high-impact initiative and applying the knowledge it generates to AI initiatives elsewhere, companies typically create more value in the new areas.

Lesson 3: Design your AI tool based on operational realities, not theory.

The day-to-day reality of how businesses operate is often quite different from how they are intended to operate. To identify what's happening on the ground, AI solution designers should shadow personnel in the business. By observing how job dispatchers and technicians use workplace tools—understanding the constraints they face and seeing how they carry out their tasks—designers can deliver initiatives that are effective in the real world. Furthermore, by applying the lessons learned to their initiatives, designers can ensure their solutions are adopted by employees rather than sitting unused and unloved on the shelf. (See "Case Study: Increasing Workforce Productivity.")

– Case Study: Increasing Workforce Productivity

A European integrated energy player with a sizeable renewables portfolio and distribution grid business faced a labor problem in its frontline operations. While the volume of tasks was increasing, the hiring and onboarding of technicians to handle the work could not keep up. Company leaders asked BCG to deliver an AI solution that would have a rapid and significant impact on worker productivity.

Following an initial assessment, the company selected a dozen use cases involving measurable efficiency levers. Among these, one of the biggest opportunities for AI to improve productivity was task execution in the field. In this area, dispatchers assign work orders to technicians, who then carry them out while facing multiple challenges including fixed customer appointments and last-minute disruptions.

The business case was framed from the beginning in terms the CFO could understand—as a means of creating value for the company by reducing technician "gap time" (the period when a technician is available but not executing work orders), increasing overall worker productivity, and avoiding the hiring of additional employees.

The project team started by selecting a small number of measurable KPIs that would be impacted by the planned solution. In addition to technician gap time, they included the time spent by dispatchers reassigning jobs following real-world disruptions, work order throughput, and avoided extra headcount.

Instead of aiming to deliver a single big number, the team used different scenarios to assess the initiative's financial impact. Observing that the average daily gap time per technician was normally between 1 and 2.5 hours, it modeled the outcome from eliminating a few hours of gap time. The team then worked with the company to achieve its gap reduction goal, lowering operating costs by a few million euros.

Rather than re-engineering the company's entire maintenance processes, our team concentrated on easing the constraints holding back worker productivity—namely, the quality of planning worker schedules and task allocation in the field and the ability to replan quickly following disruptions.

The project team found that dispatchers lacked visibility about technicians' whereabouts and their progress with carrying out jobs, and they had no unified source of knowledge about the technical requirements of different assets. Consequently, they overestimated the time needed to perform a task and failed to use gaps in technicians' availability for other work.

Based on these challenges, the team set out two key goals for the solution:

- **An upfront reduction in gap time** through more realistic planning (better estimates of task duration, improved equipment and skills checks, and smoother work order assignment).
- **Addressing any remaining gap time dynamically** through effective replanning when disruptions or early completions occur.

To create a design reflecting day-to-day realities, the BCG team spent time observing dispatchers and technicians as they went about their jobs. This brought to light several important truths and made it easier for the team to eliminate weak design ideas and focus on those supported by operational reality:

- Many instances of gap time were not the result of technicians' work habits. They were created by hard constraints such as inflexible appointments or safety requirements.
- While slack in the system was the result of inefficiencies, periods of downtime helped protect workers' well-being. This finding provided the team with a valuable lesson: the solution should not aim to eliminate every

minute of slack, but to separate healthy buffer periods from avoidable waste.

- Dispatchers and technicians used a mix of systems and temporary workarounds to communicate and solve day-to-day problems. Even at critical moments, dispatchers still used traditional communication methods such as phone calls and emails to establish technicians' whereabouts and their availability to carry out tasks.

Working within these realities, the team was able to produce an AI-enabled dispatching assistant that optimized daily planning activities and technician deployment based on well-organized information, tracked technicians' job status via push notifications, and flagged risks. The solution was integrated into real-world workflows at the MVP stage and designed to be responsive to actual behaviors while also protecting technicians' safety and well-being, enabling a successful rollout across the organization.

Lesson 4: Move fast—but with disciplined governance.

Given their role in national energy systems, renewable energy companies cannot follow the Silicon Valley ethos to "move fast and break things." But they also cannot afford to spend a lot of time tinkering with AI pilots, which can slow momentum. Business cases can change quickly. Speed is always of the essence, but it needs to be achieved responsibly.

If players are not actively using pilots to test and kill weak ideas, they are not learning fast enough. But if they are risking operational stability with a flawed solution, then they are going too fast. The right balance comes down to effective governance. Companies can deploy multiple mechanisms to achieve this balance:

- Rank potential use cases at the outset, based on risk exposure and value creation rather than the enthusiasm of senior managers.
- Decide where in the business to apply AI, using the top use cases as a guide.
- Start small and expand outwards, introducing initiatives to areas that are increasingly critical to the business.

- Establish human-in-the-loop guardrails; set non-negotiable rules, safety envelopes, and stop gates that limit AI actions; and ensure that safety compliance is reflected in engineering documentation.
- Introduce rollback mechanisms that take the initiative to a prior, reliable state when a bad update occurs, and design for graceful degradation so that it continues to operate at reduced capacity in the event of a wider system failure.

Lesson 5: Build your minimum viable product into real-world workflows.

Companies shouldn't treat their minimum viable product (MVP) as an isolated sandbox demo with no impact on real-world operations. They should view the MVP as the smallest version of an AI solution that has the potential to change behaviors at the start of the work week. The success of a solution rests upon its adoption in the workplace. If employees can ignore the solution without losing out on a key benefit, then the MVP is not doing its job.

To achieve this goal, MVPs need to meet certain criteria. For starters, they must be embedded in real-world workflows and processes. But this can be harder than it sounds. Workers may be reluctant to give up legacy systems or tools to make way for the new solution. Alongside effective communications and change management, companies may also need to implement a forced decommissioning of legacy systems to avoid new and old ways of working continuing to operate side-by-side.

To improve the chances of workplace success, MVPs must be owned by users in the business and provide support to facilitate adoption. Furthermore, alterations to the operating model will likely be needed to create new reporting lines and changes in roles and responsibilities.

Lesson 6: Scale through ownership, cadence, and value tracking.

Scaling AI initiatives from the initial development stage through rollout may look like a simple technology matter. But while technology is important, scaling ultimately comes down to operating model factors—having the right processes and people in place to achieve success.

Our work has shown that value creation through scaling cannot be achieved by an innovation team. Any group given this critical task needs to have a direct line to the business and to other key areas. Consequently, companies must establish a cross-functional unit made up of representatives from the business, IT, product development, and change management—and ensure that members are properly accountable so that they own their decisions.

The most effective operating model is a hub-and-spoke arrangement that combines centralized standard-setting with decentralized execution. At the center, a steering committee oversees AI strategy and makes investment and prioritization decisions—while a hub of senior individuals (organized around change management, HR, AI portfolio steering, the technology stack, data, and responsible AI and risk) defines group-wide standards and acts as a sparring partner.

Decentralized business units (the spokes) retain ownership of use cases and implementation within their areas, ensuring proximity to the business while leveraging shared infrastructure and governance. A time-limited centralized program can be layered on top as an accelerator to actively push priority initiatives, complementing the underlying operating model.

Building momentum and an effective cadence into the scaling process is essential. Companies should avoid long pauses between different stages—such as proof of concept, MVP, and rollout—to maintain momentum and track the value created at each stage.

AI implementation in the renewable energy sector has huge potential. However, creating successful initiatives is not straightforward. We've found that without a structured approach grounded in real-world challenges and realities, renewables players can easily fail to create meaningful value from the technology. But by following the lessons that we've distilled from working closely on AI projects with our clients, companies can harness AI's power to dramatically improve efficiency and transform their businesses.

Authors

Christopher Tuot

Managing Director & Partner
Düsseldorf

Anita Oh

Managing Director & Partner
Sydney

Frank Klose

Managing Director & Senior Partner
Düsseldorf

Andrej Levin

Managing Director & Senior Partner
Hamburg

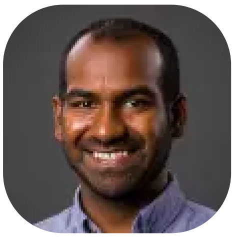

Vinoj Pillai

Partner
San Francisco - Bay Area

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.