McKinsey & Company

The AI-enabled utility: Rewiring to win in the energy transition

AI and other digital technologies can help make processes faster and more efficient, incorporate more data into decision making, and lead to higher-quality outcomes internally and externally.

Compendium
September 2023

Contents

Contents (continued)

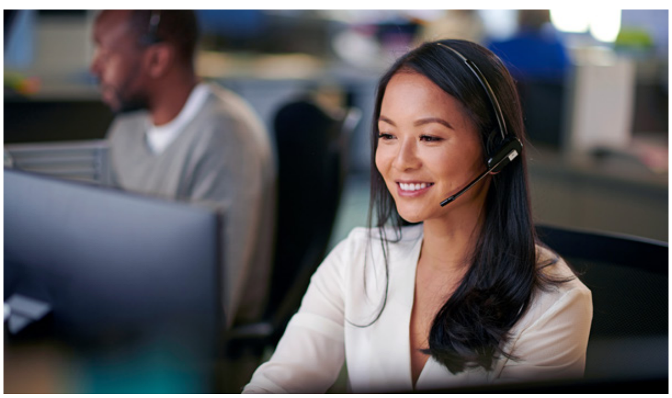

Introduction

Managing the energy transition is top of mind for today's utility and energy leaders. The challenges associated with it are manifold: utility and energy companies will need to manage decentralized power generation and demands for decarbonization while meeting rising expectations for customer service. Across all these areas, AI and other digital technologies offer powerful solutions, representing a major opportunity to overcome the challenges of the energy transition and enable future success.

However, AI and other digital technologies can be difficult to implement, and utilities have historically struggled to create value from them. Legacy systems and decentralized organization structures can lead many digital transformations to fizzle out before they can take effect. As a result, many players have failed to reap the benefits of technologies such as machine learning and advanced analytics, leaving them less efficient and less resilient in the face of external pressures.

This compendium offers numerous perspectives on and examples of digital transformations in utilities. These include deploying digital platforms on top of available technology foundations, leveraging advanced analytics to elevate the customer experience, implementing advanced analytics in asset management, and leveraging AI-driven schedule optimizers to improve productivity.

If utility and energy leaders can effectively digitally transform their companies, they have the potential to capture efficiencies beyond those of their competitors. We hope this compendium offers new insights to help them succeed on their journeys.

Digital transformation in energy: Achieving escape velocity

The need for digital value is greater and the barriers to change are lower—yet inertia persists. Three practical lessons can help energy companies break through.

by Adrian Booth, Nikhil Patel, and Micah Smith

For energy companies, achieving value from digital technologies has become the great white whale: anxiously hunted, dimly perceived, enormous, and elusive. For the past two years, energy companies of every stripe have probed digital possibilities—running pilots in analytics, process digitization, and automation. They assumed that, as engineering-savvy organizations with a history of ingenuity, they could easily find the value from digital. Reality has proved more difficult.

Energy companies have failed to achieve substantial business value from digital because their approaches do not account for the unique challenges of being an energy company, which create extraordinary inertia. Breaking that inertia will require far bolder action than energy companies have been comfortable taking to date. They must commit to transformation.

To be sure, COVID-19 has scrambled the outlook for digital in energy. Some companies, especially independent producers and suppliers in oil and gas, must focus urgently on cash and survival—digital can come later. But for energy companies with the resources to weather the storm, the disruption of COVID-19 has done two things: first, it has underlined that survival requires getting to the next level on cost and adaptability, and that requires digital; and second, by forcing companies to abandon business as usual, it has lowered the barriers to change that typically impede digital transformation.

Digital transformations have been tragically overhyped, but we believe they are both possible and necessary. In this article, we describe a perspective on what it will take to achieve one in energy. We draw on lessons from early experiences in energy as well as earlier movers in high tech, finance, healthcare, and mining. These lessons were true before COVID-19 and we believe they remain true now. Energy companies that harness them will have the wherewithal to invest for future leadership.

Why enabling digital in energy matters

So far, the adoption of new technologies in energy has been more hype than reality. Does it warrant so much attention?

It is truly important that energy companies realize the promise of digital innovation at scale, on a global basis. It matters to the world: over the next two to three decades, more than five billion people across the developing world will seek a path out of poverty. Unlocking the magnitude of energy resources required to improve their lives, in a way that does not choke the environment, cannot be done without the power of digital to improve efficiency and manage complexity.

And it matters to energy companies because they face unprecedented changes across the energy system: more competition, more complexity, and less predictability. Profit margins are under pressure, and the margin of error for survival is shrinking.

These changes affect every player:

— *oil and gas operators* that face price volatility, potential peak demand, and the dynamism of shale versus OPEC

— *utilities* that face distributed generation, more complex grids, and evolving customer expectations

— *refineries* that must adapt to global uncertainty over new sources of feedstock and new patterns of demand

— *renewables developers* that must survive and grow amid intensifying competition and potential commoditization

— *service companies* that must remake their delivery models to meet customers' new expectations about digital efficiencies

— *engineering, procurement, and construction companies* that struggle to deliver the types of capital projects that matter for the future

Digital innovation is one of the few means that can contend with these profound changes—by using predictive analytics to better anticipate the future, data to better inform current decisions in the here and now, and digitization and automation to take advantage of every increment of cost and speed that can be found.

In the narrow settings where energy companies have applied digital successfully, we have already seen it facilitate 2 to 10 percent improvements in production and yield and 10 to 30 percent improvements in cost. If these benefits hold true at scale, they could have a material impact on competitiveness: for example, improving cost efficiency by one to four cents per kilowatt-hour in power and $2 to $12 per barrel in upstream oil and gas production (Exhibit 1). These deltas can make the difference for a business's survival, let alone its leadership.

Challenges of digital in energy

Why have efforts to gain value from digital not succeeded? In fairness, energy companies have learned many lessons from their frustrations over the past two years. These lessons are now so widely repeated that they have almost become clichés:

— Focus on the business case, not the technology.

— Listen to and delight the users.

— People, process, and behavioral change are just as important as technology.

— Avoid pilot purgatory.

— Do not let the perfect be the enemy of the good; instead, fail fast and learn fast.

— Agility is good; burdensome hierarchy is bad.

— Transformation takes effort, whether digital or not.

These lessons are true but insufficient.

Energy companies sometimes overuse being different as a crutch to explain why innovations from other industries cannot apply to them. That said, they do face different circumstances from the digital-native tech pioneers of Silicon Valley or the financial institutions that led the next wave of digital and had business models that entailed largely nonphysical, easily scaled transactions.

Exhibit 1

Done right and at scale, digital will materially affect competitiveness.

Narrow efforts at digital have shown ~2 to 10 percent yield improvements and ~10 to 30 percent cost improvements in capital, supply chain, and operations. **What is the cost efficiency opportunity if these impacts hold at scale?**

Illustrations of potential impact

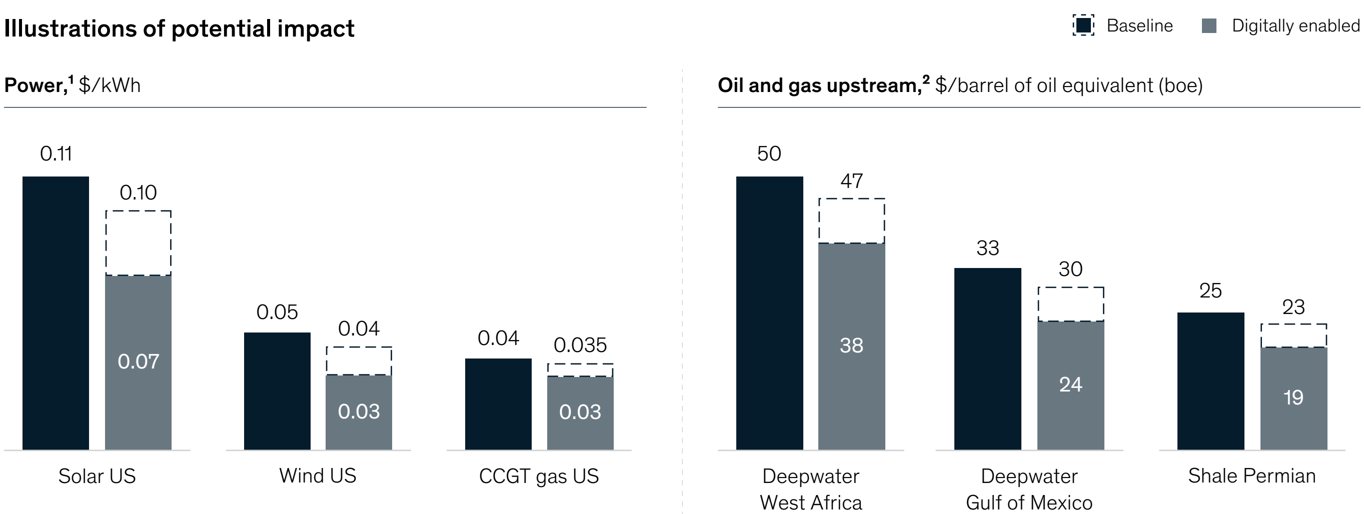

[1] Assumes US average figures for all operating plants, and midrange capex from listed sources for new build.
[2] Per boe cost estimates based on assets that came online 2010-19; opex and government take based on 2019 spend and production; capex per boe for deepwater based on total investment spend from start of development to two years after start-up divided by total resources; capex per boe for Permian is based on well capex divided by estimated EUR; deepwater assets defined as having water depth of more than 1,000 feet.

Source: *2018 wind technologies market report*, Department of Energy, 2018; Energy Information Administration; "Lazard's levelized cost of energy analysis, version 13.0," Lazard, November 2019; Rystad Energy UCube; *Utility-scale solar, 2019 edition*, Berkeley, December 2019; McKinsey analysis

We have already seen digital facilitate 2 to 10 percent improvements in production and yield and 10 to 30 percent improvements in cost.

What makes digital value difficult to achieve in the energy industry?

— *Physical orientation.* The energy business is sensitive to the laws of physics—whether the geophysics of an oil and gas reservoir, quantum physics of solar power, fluid dynamics of wind, thermodynamics of fossil power, or electromagnetics of power transmission. Moreover, it is embodied in heavy capital such as power plants, offshore platforms, or liquefied natural gas terminals or pipelines. This physicality makes energy operations, and therefore profit generation, fundamentally difficult. In energy, digital applications must contend with the laws of nature and be done in a way that safeguards asset health and frontline capability. And proposed technology investments must meet a high bar of proof that they are worthy of integrating into these difficult operations.

— *Health and safety risk.* Energy is a powerful commodity: it supports our everyday lives, but without care, it is potentially dangerous. The industry pays enormous attention to safety, but incidents still occur—sometimes deadly, such as the Deepwater Horizon and San Bruno explosions. Given the inherent risks, energy companies must navigate a web of regulations, which span every level of government. For example, there are local rules about land use; state rules about water, safety, energy mix, or consumer service requirements; federal rules about interstate projects and operations; and even international treaties governing the energy trade. As a result of living under regulatory scrutiny for more than a century, energy companies are averse to risk and try to control for it through detailed and rigorous processes. This makes them slow to change.

— *Engineer-driven culture.* In a physically oriented, highly regulated sector, the engineer is king. Oil and gas and power companies are dominated by current and former engineers who have risen to the executive ranks. As such, these organizations are enmeshed in an engineering mindset: a love of large, comprehensive projects, a premium on finding the perfect solution up front, detailed planning to the highest degree, and a preference for rigorous analysis and process over fast judgments and flexibility.

— *Heavy dependence on third parties.* The work of energy companies depends on an extensive and fragmented supply chain. Industries such as airlines and automotive also depend on a complex supply base, but the energy industry puts supplier collaboration at the heart of operations. For example, a shale play requires a huge cascade of parties—owner companies that provide funding and regulatory engagement, third parties that drill, others that haul water and sand, others that build surface facilities, and others that integrate all these activities—just to produce a barrel of oil.

— *Long careers, narrow exposure.* Many energy executives have been at the same company for at least 30 years (unusual for today), rising through the ranks by running a well-worn playbook. Such discipline makes sense given the complexities of managing a business that is physical, high risk, highly engineered, and fragmented. These people were rewarded not on innovation but on caution and following tradition and are thus more practiced at surviving business cycles than driving sustained change.

— *Global operations.* Energy companies go where the resources are. Often, this requires companies to span remarkably diverse regions: from Texas to Angola, California to India, Germany to Indonesia, the Persian Gulf to China. The legal and operating environments in these regions vary dramatically; even relatively simple things such as internet connectivity cannot be taken for granted. Some governments are content to let markets be, while others use energy investments to promote national development. Labor forces vary in capability, reliability, size, and cost. Supply chains vary in maturity. And legal recourse runs the gamut from impartial courts to relying on knowing the right people.

Other industries contend with one or two of the issues above; the energy industry deals with them all. Each factor encumbers movement. Taken together, these challenges create massive inertia when it comes to digital adoption.

How to break the inertia: Digital fortune favors the bold

The only way to overcome an organization's inertia is to apply enough force. Small steps such as pilots and proofs of concept are too weak, and large technology programs run the risk of creating the wrong kind of force. The business must commit to digital transformation—fundamentally changing how the organization works, beyond technology alone—and then go after it.

COVID-19 has created a window to drive transformative change, in small part because companies and workers have been forced to accelerate their use of basic digital tools for remote work, and in large part because the status quo has already been disrupted.

It is easy to say "transform," but what makes a digital transformation stick?

After much experience in the trenches, we have developed a digital transformation journey that breaks the inertia, unlocks large-scale value, and lasts (Exhibit 2).

Our digital transformation journey reflects three practical principles to tackle inertia and achieve scale:

1. Don't just sponsor—own.
2. Don't just create tools—transform whole workflows.
3. Don't execute the transformation in one big bang—take bite-size actions.

These three principles animate every step of the digital transformation journey. How are they translated into practice?

Don't just sponsor—own

Most energy companies have a matrix of business units, which run the operation, and functions (such as IT, engineering, or procurement), which provide services to the businesses. Usually, businesses, not functions, have the power. Digital transformation can happen only if businesses take responsibility for transforming themselves.

Almost invariably, they don't. The mindset of business executives has evolved over the years, but it's still not enough. Two years ago, business leaders relegated digital to the sidelines as a quirky IT project. Today, they recognize the value of digital but still treat it as a visionary experiment from on high. Indeed, corporate leaders hail the promise of digital technology and fund it lavishly, but they make it the purview of new digital incubators or centers of excellence rather than line operations.

Business units will not commit to digital transformation unless they have skin in the game. What does that look like?

Exhibit 2

A digital transformation journey in energy requires reimagining workflows.

Value-focused vision within weeks; value and action within first 6 months; transformation in 18–24 months

	Road map	Vision (by workflow)	MVP[1]	Industrialize	Scale and expand	Platform
Value unlock	**Define the end-to-end workflows** that drive the most value in the business ("needle movers")	**Reimagine future workflows** to get the most value	**Rapidly deploy initial products** to users to deliver value fast, generate learnings, and create a springboard	**Harden the MVPs**[1] to make sure they will work in live operation at scale	**Realize full vision** by expanding beyond MVPs,[1] reusing across business units, and building new products	**Establish a sustained digital factory** that is an engine of enabling digital for the enterprise
Data and technology	**Conduct rapid gap analysis** of tools and infrastructure	**Inventory key systems** of record and field, pilot, or planned technologies	**Institute basic best practices,** including API-first approach, rationalized tech stack across business units, and automated security approvals	**Clean the code,** enabling scale-up **Institutionalize tech enablers** (eg, site reliability engineering)	**Create code libraries** for common needs, and instrument the code to enable performance analytics	**Create an API marketplace** that makes the reusable building blocks available to all for continuous innovation
Culture and capabilities	**Conduct rapid gap analysis** of digital and nondigital capabilities	**Engage** the most courageous, informed, creative leaders to own and shape the vision	**Catalyze frontline buy-in** from business units and create a forcing mechanism to simplify IT policies	**Establish user support** process and capabilities to ensure manageable scale-up	**Demonstrate the value** of sharing, standardization, and scale **Expand in-house talent** base	**Formalize the digital factory's** operating model and replicate it
Time frame	4–12 weeks	4–12 weeks	8–12 weeks	4–6 weeks	2–8 weeks for reuse; expansions span 12–18 months	After 12–18 months

[1] Minimum viable product.

Being a sponsor and providing funds signals interest, but it does not generate a psychological sense of accountability. If all a business does is fund a digital service that another group must deliver, it can remain at arm's length as a dispassionate judge of someone else's work.

Beyond funding, business units should first integrate digital technologies into their formal business and operational targets. This has a dual benefit. On one hand, it forces the business to take accountability for creating value from digital. On the other hand, that accountability also gives the business the power to shape digital initiatives to reflect true priorities, not the fanciful dreams of technologists.

Second, business units should commit the time and talent of their star operators. Typically, business leaders who view digital innovation as a sideshow shield their stars from digital efforts and refocus them on the traditional moneymakers: drilling wells, signing power purchase agreements, operating the

COVID-19 has created a window to drive transformative change . . . because the status quo is already disrupted.

Example: Creating business accountability for digital

At one renewables company, senior leaders recognized that digital innovation would become important for the enterprise to remain competitive amid an environment of more competitors, tax credits phasing out, and commoditization. Rather than setting up a special digital program, they pushed all the business units to define their own digital initiatives, as part of their core process of strategy, planning, and target setting. Crucially, the businesses did more than come up with digital initiatives: they had to assign a value aspiration in dollars and set their budgets accordingly. Doing so ensured that digital initiatives were practical and worthy of the same priority as growing the project backlog or keeping turbines operational. The businesses did not fly blind—they enlisted help from IT and analytics groups—but whatever initiatives they conceived, they owned with conviction. At the same time, senior executives watched for sandbagging and pushed the businesses to stay ambitious.

plant most efficiently, and so on. But if a business wants to turn digital into a moneymaker, too, it needs to put its best people on the case.

Star operators know the culture of the business, know the operation, and have good judgment about how to make changes. They are the most credible champions of a digital effort since they are widely respected by even the flintiest 30-year veterans, and they also have the capability to be creative about the future. Obviously, a business cannot afford to keep its best talent focused indefinitely on digital innovation, but by cross-pollinating between digital and core operations, they will eventually become one and the same.

Third, businesses should make the funding commitment big and bold. Digital transformation depends on deep operational change—weaving technology into operators' daily lives, streamlining processes, and radically accelerating the pace of activity. Those changes are subject to fits and starts and reversals, so it is hard to predict which changes will succeed or fail and how best to adjust. To make all this work, businesses should devote a big block of funds that covers a long time and gives digital efforts room to zigzag toward transformational change.

Don't just create tools—transform whole workflows

The prevailing approach to digital innovation in energy is to create "point solutions," which apply technology to narrow uses without fundamentally changing how people work. Energy companies succumb to point solutions partly because they have delegated digital to the IT function, which operates through a technology lens and lacks a mandate to redefine how other functions do their work. Moreover, energy executives are prisoners of precedent: when they hear "digital," they conjure images of 4-D seismic data upstream, advanced process controls downstream, or grid-operations centers in power—in short, tools to digitize the existing process rather than a wholesale reimagining of the process.

Example: Block funding with performance accountability

A major oil and gas company adopted a practical "block funding" approach that balances flexibility with accountability. The company approved an annual budget for its digital transformation, with quarterly check-ins with the business leaders whose business units are key participants in the transformation. The budget is large enough to empower investments in big changes. Between check-ins, the business units are free to innovate, fail, and adjust however they deem best. In the quarterly check-ins, they review quantitative estimates of value (such as reduced cost or improved speed) and qualitative proof points of value (such as user adoption rates or testimonials from the field) and jointly decide whether to stop, adjust, or proceed without change. Learnings from this cycle are woven into the next block fund.

Point solutions also fall prey to the perils of inertia: they are not meaty enough to change behaviors or mindsets, they often digitize poor processes rather than improve them, they get smothered by the status quo, and they optimize targeted needs instead of the whole system.

To break inertia, digital transformation requires a forceful reimagining of how people work—and the only way to do that is by rewiring entire end-to-end workflows (Exhibit 3). Workflows can be defined in a variety of ways, but they should reflect the main vectors by which a business generates value. Examples include production optimization in upstream oil and gas, real-time turnarounds in refineries or petrochemicals, and plant operations and field maintenance in power.

Focusing on end-to-end workflows is vital for success in digital transformation. First, it launches a deep rethinking of how an operation works, which generates the creativity and momentum required for true transformation and points everyone toward

Exhibit 3

End-to-end workflows have value greater than the sum of their parts.

Maintenance optimization example

Workflow segments	**Determine maintenance needs**	**Dispatch and route crews**	**Execute maintenance**	**Log and track**
Challenges	• Preventing asset failures • Knowing where and when failure might happen	• Low crew utilization • Unnecessary trips • Knowing where and how to direct crews	• Rework • Jobs take longer than needed • Quality variance	• Limited, poor-quality data • Manual entry takes time away from work, introduces error
Point solutions alone *Useful but limited*	**Predictive maintenance algorithm** identifies assets nearing failure	**Remote restart** attempts to solve issue automatically **Routing optimization algorithm** matches the right crews with the right place, routes them efficiently	**Mobility tools** provide crews with live access to data and remote support	**Sensor records and easy manual entry** ensure accurate record of work done and asset status

Solutions connected across a workflow
Powerful

- Predictive algorithm **informs** how route optimizer sets priorities
- Routing optimizer **feeds** into intuitive user interface so crews easily know where to be, when, why, and with what tools
- Site-level data **seamlessly feeds back** to the center to create better analysis
- Predictive algorithm **feeds diagnostics** into mobility tool

a shared goal. Second, it makes all digital efforts subservient to an ambitious business vision. To transform a workflow, a business is forced to define a vision of the future state: how the activities in that workflow can be done the fastest, safest, and most cost-effectively, and what needs to change from today. Digital technology enters the picture only as an enabler to help make the vision a reality.

With this approach, businesses find it easier to define digital initiatives that actually lead to transformation, not just incremental gains. Moreover, they can easily prioritize digital investments: out of the 1,000 ideas their organizations dream up, only ten really contribute to the future state vision—the other 990 get dropped. Finally, they can more easily rally the organization around the effort, because the benefits of digital become more obvious—even to skeptics.

Example: Making technology relevant to how people work

One US power company thought there might be value in drones. It started with a "technology-back" approach: It bought a small fleet of drones, which each cost around $3,000, and then looked for applications (the most obvious being inspections).

At first, the company used the drones for external inspections of boilers: an operator on the ground, with line of sight to the drone, flew it around the boiler to capture video feed. This was promising, but relatively low value.

It soon became apparent that there would be more value in sending the drones inside the boilers while they were still hot, which would allow the company to run inspections without having to shut down operations.

But because this company had focused on technology first, it had not fully considered the workflow.

Operators needed a way to inspect the boilers quickly and thoroughly without any risk of damaging them. The original "user journey" entailed shutting down the boilers and having human beings go inside to perform quick inspections (to allow operations to come back online) and comprehensively (to ensure equipment integrity and absolute compliance with regulation). Any revised "user journey" utilizing drones would need to enable this same speed and comprehensiveness.

But the company had chosen its drone technology before defining what the future workflow should be. Their drones required a human operator, who had to navigate the drone within the boiler. At first, the operators were excited about the drones, but they soon became fearful about collisions within the boiler—so they simply stopped using the drones.

Leadership decided to recalibrate, directing their attention to the workflow. To get the value of internal boiler inspections, the workflow demanded a technology that could work quickly and thoroughly and provide confidence to operators. With that in mind, they went back to market and found a start-up with a drone that could fly itself, using machine vision to avoid obstacles (at one-third of the price of the original drones, no less). Moreover, they considered the full opportunity for the future workflow: What if the drone not only provided video feeds but also used machine learning to diagnose what it was seeing to identify leaks? The company is now on a path toward a much higher value ambition, with much higher adoption by the organization.

Don't execute the transformation in one big bang—take bite-size actions

Inertia is tenacious—it will resurface over and over. Energy companies routinely try to plan every aspect of their digital efforts up front. Despite lip service to agility and modularity, they aim for engineered perfection. All of that planning and perfection leads to helplessness when momentum lags or unforeseen problems emerge during a transformation.

With digital technology, the right answer is hard to lock in right away. For one thing, technology changes at a rapid clip—current trends could become obsolete in less than a year. Furthermore, it's impossible to know what will work until it's underway—building something real, giving it to users, and testing it against reality.

Therefore, a digital transformation is best advanced in bites. Select a digital solution, cocreate it with field users, get it into those users' hands quickly, and let them start using it to generate value—and, by the same token, let them get excited by it. Learn from the experience, and play those lessons into the next push. If a need to adjust arises, multiple pivot points are available. This approach offers repeated surges of action and progress to continuously break inertia.

The numerous merits of the bite-size method have been validated many times over by energy companies that have adopted it.

One benefit is that it delivers business value and creates supporters quickly, which generates self-perpetuating momentum for the digital transformation. And as we have seen, momentum is everything.

Example: Creating momentum through surges of value and progress

A leading oil and gas company has undertaken the bite-size approach to achieve unprecedented value, speed, and scale of transformation in end-to-end procurement and supply chain. It outlined an ambitious vision for the future of procurement and supply chain and took a starting action quickly by launching three digital products in parallel, in three business units. The company built the products and put them in users' hands within two months (including tough-to-implement elements, such as predictive analytics used by operators in the field), and scaled between businesses and asset types (which were as wide-ranging as shale, deepwater, and conventional onshore) in four to six months. The businesses were excited to reduce operating costs and speed up cycle times, and the frontline operators were thrilled that their lives were made easier. Word spread—creating pull from the rest of the enterprise.

A second benefit of the bite-size approach is that it forces practicality. Time pressure focuses the mind on what matters.

Example: Pragmatic compromises and being quick to action can lead to better outcomes

A North American shale operator sought to transform its water logistics workflow, supported by digital technology. But making it happen would be tricky: it needed to influence a long tail of third-party water haulers of varying capability (over which it had limited control) and to create visibility across its network of hauling operations to improve cost and safety—especially important given transit through populated areas.

The answer most likely to be pitched to executives was to push all haulers to adopt standardized sensing devices selected by the company; to create instant data feeds of where every truck is at every second; and to create sophisticated algorithms to optimize truck movements and direct them centrally. But they did not have time to gestate these grand notions—they wanted to take action. And after analyzing the workflow, they realized they did not need perfection: they could capture the bulk of the business potential through bite-size variants. Instead of pushing standardized sensors, they made use of the devices already installed by different hauling companies on their respective trucks, and they found simple methods to reconcile data from disparate systems. Instead of aiming for real-time network visibility and optimization, they realized that a “day ahead” view would do the trick.

Finally, there is simply no substitute for experience. People often struggle to understand the full value of a digital initiative on paper, no matter how well described or planned. It’s only once they see the reality that they grasp the potential. In one case, a shale business unit built a demand-planning solution for drilling logistics using advanced statistical models to predict when material orders would be needed based on the depth of well, the speed of drilling, and other factors. One executive involved in the effort was supportive but believed the solution was limited to its origin (drilling logistics in shale). Over the course of two months, however, her outlook shifted. She spoke not only with users in the field, seeing how they actually used the predictive solution, but also with operators in other businesses to understand their needs. Over time, she appreciated just how big the opportunity could be—not just predictions for drilling logistics, but for all logistics; not just materials, but also services; not just this shale asset, but every asset.

A couple of concerns about the bite-size approach are worth considering. The first is that it leads to one-off solutions that cannot scale. In our digital transformation journey, reusability and scalability both factor into decisions from day zero, which are about defining the future-state vision and about delivering the first minimum viable products. When done properly, we have seen MVPs scale up to new business units and masses of new users with 80 to 90 percent of the original code unchanged.

That said, the concern over scalability may be a red herring. A digital program can be regeared to scale, even if scalability was completely ignored in the past. For example, a leading financial-services company had worked to transform multiple workflows but had proceeded in silos. Each team invented its own approach to user access, business rules, data integration, and so on. As a result, a user who started in one workflow could not seamlessly move into another workflow—for instance, after a customer onboarded to a bank and received a loan, she could not automatically open a debit account. As the user base grew, this element became a real challenge. But to achieve interoperability and unlock scale, the company did not have to suffer months of agony redoing its code. Instead, it simply took a two-week “step back” to revise the workflows to act as an interconnected whole, identify the reusable components that were common across workflows, and set up the integration patterns to allow for API enablement of data flow between systems. All it took was two focused weeks to reposition the effort for scalability, and it provided a huge accelerant. The company was able to go from two workflows to four to eight in parallel in about six months, and it improved efficiency by more than 30 percent through automation and reuse.

A second concern with the bite-size approach is that it may work for digital product development, but not for establishing technology foundations, developing capability, or changing culture, which require more planning and longer time frames. This notion has been repeatedly disproved by experience. When it comes to technology, energy companies have tried laying all their foundations up front (enter the dreaded "data lake"), but nobody races to use them. When it comes to capability, they have hired a gaggle of pedigreed data scientists and said, "Go forth, mine the data, find the money," but most of those new hires sit around waiting for a purpose. When it comes to culture, they have run interminable workshops but never get beyond classroom theory.

By contrast, the bite-size approach tries to advance the trifecta of business value, technology platforms, and culture and capabilities through each step of the transformation—not in theory, but through delivering on a specific business need. Any time a technology choice or a hiring push is made, it is made because it is necessary to deliver on an immediate business need. Cultural changes come from actual changes to how people work and interact with each other in the field.

For now, digital transformations in energy will largely focus on operations. That scope is hard enough and has plenty that needs to be addressed. But in successfully reimagining operations—and building digital capabilities along the way—energy companies will open the next horizon of digital opportunity: truly disruptive business models. We are at only the beginning of the journey.

Energy companies are right to take digital seriously; it is important for their future success, and it is imperative for global economic growth and environmental care. But they must summon the boldness that marked their triumphs in the past, making ingenious innovations a reality—whether creating a continent-spanning grid, extracting hydrocarbons safely from the planet's most remote areas, or using quantum mechanics to harness the energy of the sun. As in those endeavors, so with digital: fortune favors the bold.

Adrian Booth is a senior partner in McKinsey's San Francisco office, **Nikhil Patel** is a partner in the Houston office, and **Micah Smith** is a senior partner in the Dallas office.

The authors wish to thank Claudio Brasca, Collin Cole, Eelco de Jong, Travis Dziubla, Brennan Gudmundson, Spencer Holmes, Bharat Kejriwal, Ling Lau, Murali Natarajan, Thomas Newton, Jesse Noffsinger, Craig Poeppelman, Abhay Prasanna, Michael Rix, Naveen Sastry, Lois Schonberger, Thomas Seitz, Chhavi Sharma, Anjan Siddharth, Meera Chidambaram Sivakumar, Kyler Tebbutt, and Liangliang Zhang for their contributions to this article.

Winner takes all? Digital in the utility industry

With the utility industry fragmented across North America, can digital platforms enable it to consolidate at scale, resulting in lower prices, better service, and more satisfied customers?

by Adrian Booth, Eelco de Jong, Ben Elder, and Aditya Pande

Load growth is slow, energy prices are soaring, inflation is rising, and grid reliability and resilience is becoming an ever-present concern—North America's population is under pressure and would welcome an easing on their wallets from the utility industry. Meanwhile, most utilities' bold aspirations to reduce their carbon footprint to net zero over the next few decades are being met with capital, labor, and materials challenges that make achieving this goal uncertain. For public utilities, these challenges are only further amplified by earnings pressure amid a volatile energy market. And the grid is only getting more complex to operate as distributed energy resources introduce an influx of new information and variables into the system.

Key questions to be asked

With these issues in mind, solutions need to be sought. The North American utility industry is massively fragmented—more than 3,000 electric utilities and many more gas and water ones—across three primary ownership models that are either investor- or municipal-owned utilities or cooperatives.[1] This raises questions:

Why hasn't the industry consolidated more to take advantage of scale and best practices to deliver a better product at a lower price with higher customer satisfaction? At this stage, it is very difficult to prove that the industry can provide this. Many proposed M&A strategies run into well-meaning, state-based utility commissioners, city managers, or cooperative shareholders who may wish to protect local communities and potentially disallow the typical M&A deal synergies—such as reducing corporate overhead costs and operational expenses that could risk service levels or other actions, leading to layoffs, higher customer rates, or reduced service levels. When trying to compare utility performance, discussions often quickly get lost in the nuanced differences of each utility—such as whether it is urban, suburban, or rural and overhead or underground; weather and vegetation variations; historical capital-spend levels per customer; or age of assets.

What if there were a way to build a utility that could demonstrably prove that it could provide a better product and service at lower rates? Digital disrupting business models across industries are increasing rapidly. This began with entirely new sectors being created (such as search or social media); some of the companies in those sectors are now among the world's most valuable. Then disruption moved to mostly "asset light" industries: those in which the product or service was primarily based on information or data such as banking, media, or insurance. The disruption then traveled to industries where physical products were involved, such as e-commerce. Now the disruption is increasingly blurring the lines between physical and digital, such as Tesla and Peloton Interactive, whose combined digitally infused physical products are fundamentally superior to alternatives. Despite this, the utility industry barely takes advantage of digital.

What if monopoly-based sectors could use digital to disrupt the monopoly structure? Because digital has not yet fully infiltrated the utility industries, what would happen if the regulated utility networks of electric, gas, and water businesses could use digital to deliver electrons and gas or water molecules in an alternative fashion? At present, there is little evidence that the industry is pursuing such innovation at the same scale and pace seen in other industries, so why not flip the question on its head and ask: If digital could enable a utility to provide a fundamentally better product and services at lower rates, what could that do to the utility industry's underlying structure?

In this article, we explore answers to these questions, expose the significant opportunities that the space presents, look at what is needed to build a digital utility platform, and identify six key factor for success.

A modern digital platform: A once-in-a-generation consolidation opportunity

Digital could open an exciting consolidation opportunity for fast-moving companies that create digital platforms to meet their customers' needs.

[1] *United States Electricity Industry Primer*, Office of Electricity Delivery and Energy Reliability, US Department of Energy, July 2015.

The pace of technological change is increasing—look at the fast-decreasing cost of cloud computing, the growing availability of powerful machine-learning (ML) and AI capabilities, the rapidly evolving tools to deal with persistent and chronic data issues, and the increased intelligence in smartphones.

While more utilities are starting to adopt many of these digital trends, the rate of adoption is not keeping up with the pace of innovation. The opportunity gap to improve key outcomes by deploying technologies and methodologies that have been utilized in successful transformations increases every day.

The evidence for digital is clear. When working with leading utilities, we have seen exceptional step-change improvements in select use cases such as the following:

— a 25 to 30 percent field productivity improvement from AI-powered scheduling

— up to an 80 percent capital reallocation based on ML insights in asset health

— a more than a 30 percent improvement on customer satisfaction in select journeys

— a 2 to 5 percent increase in heat rate or yield for fossil as well as renewable generation assets

— a more than 30 percent improvement in reliability and resilience outcomes within existing spend levels

If the "product" is defined as clean, reliable, resilient, safe, easy to do business with, and an affordable energy or water service, then a step change in every dimension is possible. This can be done by looking at a collection of already-achieved impacts from utilities using select digital use cases.

A digital platform on top of an available technology foundation

What would happen if a digital platform that deploys every known high-impact use case to its full extent was built on top of a flexible, extendable, available technology foundation that could "bolt on" additional utilities?

While almost all major utilities are using digital, data, and analytics in some fashion, it appears that few executive teams can articulate a cohesive strategy on how a comprehensive digital, data, and analytics platform could provide "best in class" outcomes across reliability, safety, resilience, affordability, and customer experience—with no trade-offs.

In our perspective, if a cohesive strategy is not devised within a three- to five-year time frame, likely no one will "break out of the pack" and the industry will continue on its linear improvement trajectory.

However, bold industry companies that adopt a digital platform could achieve a step-change performance ahead of peers and, more important, use the once-in-a-generation opportunity to fundamentally restructure the entire industry. The value at stake is massive for those that take action. They could invite energy regulators, customers, and communities to join an unbeatable deal—a digital utility platform that provides the best reliability, safety, resilience, affordability, and customer experience. Those jurisdictions and utilities that connect to the platform could be set up to tackle the energy transition from a position of strength. If the core utility product can be offered at lower cost and better customer experience, it will create more headroom to invest in carbon-free technologies or improvements in grid resilience, or both.

Many stakeholders would need to be involved to achieve this, including customers, investors, policy makers, and regulators. The regulatory relationship would be critical, given the authority that regulators generally have in approving (or disallowing) investments. For the vision to succeed, it would require an open utility and regulator relationship, with both willing to explore a new partnership based on transparency and verifiable outcomes.

What it will take

Building a base that can serve as a comprehensive multi-utility platform will require a detailed, layer-based framework and associated design elements.

These layers can collectively transform legacy utility architecture into a "digital-native style," secure, and interoperable platform that allows business services to be scaled across utilities (exhibit).

How to build it: A new approach and new leadership

Most utilities are already building parts of these features across some layers of their tech stack; in other words, these features are being built on a use-case-by-use-case scenario—think of it as "drip irrigating" a farm with new elements. While these enable specific use cases, the escape velocity that is needed for all the layers to be in place will take too long to deliver an efficient model quickly. Moreover, this is an optimistic outcome that will require a lengthy history of delivering cross-business use cases, alongside a visionary enterprise architecture team that can enable the organic build-out to collectively scaffold this cross-layer end state.

Building the multilayer future state will require a cross-functional team and a close partnership with the enterprise architecture organization. The team needs to set the intention to focus on creating and delivering this future state while the rest of the IT organization delivers nearer-term use cases and other "must do" regulatory or systems migrations.

A platform architecture road map

Building such a platform will be a multiyear endeavor, and utilities will need the ambition and the resolve not only to embark on the journey but also to stick with the aspirational vision over a three- to five-year horizon and buttress against shifting priorities throughout. Various key enablers are necessary to have in place up front to put utilities on the path to success when making this leap.

A strong business case backed by estimates of measurable business value will be needed to develop a long-term road map and strategy and return value to the utility. The road map could include the series of strategic investments required to deliver the transformative business-enablement platform, and an estimate of the necessary foundational investments. Further, leadership will need to keep front of mind the necessity for significant new talent and partnerships to create the blueprint and execute the build-out of this platform.

New technical capabilities and skill sets will be required, including cloud engineering, DevOps, and data science. Digital skills like these will be critical to deliver the target state platform. Utilities will need to ensure that best practices around core cloud and software engineering capabilities are in place—for example, infrastructure as code, cloud FinOps, automated testing, and cybersecurity.

A well-defined set of foundational architecture principles and a lean tech governance model will be necessary to ensure that maximum value from the investment is returned to the business. This is essential to steer strategic design decisions made along the way. An architecture governance model, backed by a shared set of principles and guardrails, could drive delivery consistency through the use of acceptable patterns, streamline technical decision making, and empower delivery teams by giving them the autonomy to move at pace with agility.

Strong organizational cooperation and commitment by multiple stakeholders will be essential, beginning with the C-team and board. Because this transformation is a multiyear journey, dedication and support from top-level leaders will be important if a company is to stay the course, with frequent and consistent communication at all levels. The development of the target state platform needs to be a collective effort—it cannot be achieved in isolation. Business stakeholders could consider partnering with the whole IT function to ensure the alignment of goals and outcomes, address dependencies, and reduce risk.

Three phases of transformation

We have observed that a successful transformation to a digital utility platform takes place over three phases.

Phase 1: Developing foundational patterns for integration architecture

The first step in implementing a target state platform model is to establish foundational patterns for system integration and platform consumption. The integration layer is a vital starting point—getting it

Exhibit

Adopting a digital platform architecture model could enable a utility to forge ahead of competitors.

Components of a digital platform architecture model

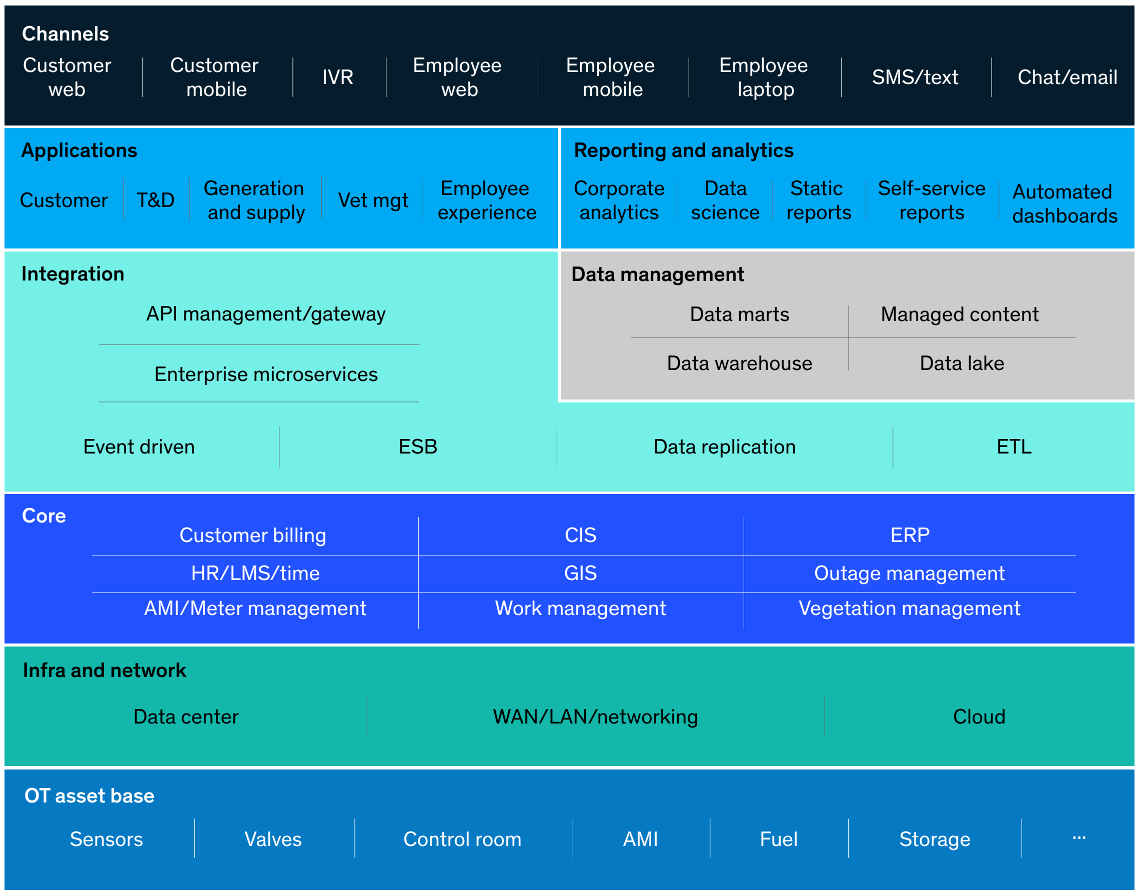

Key characteristics

- **Unified user experiences:** fragmentation of experience across apps and tools is one of the most common pain points for utility field workers today. Reusable UX components and cross-platform development help deliver consistent and seamless experiences in the field.
- **Analytics ready:** empowering citizen development through analytics sandboxes and self-serve reporting.
- **Enterprise APIs and microservices:** a well-defined catalog of domain-driven APIs enables utilities to build fit-for-purpose omnichannel solutions on top of, but decoupled from, core enterprise systems.
- **Robust core system integration:** most utilities struggle to access and leverage the data in their core enterprise systems (eg, assets, work orders, customer, etc). A modern platform will provide near real-time integration to read from and write back to core systems for analytics and app use cases.
- **Cloud native:** moving workloads to the cloud and leveraging elastic scale for storage and compute is helping utilities drive down capital and O&M costs across business units.
- **Leveraging the Internet of Things data explosion:** the abundance of data from AMI and smart sensors has largely gone untapped by utilities. A modern platform will curate and synthesize this data for predictive modeling use cases.

McKinsey & Company

right means fewer headaches during large system upgrades or consolidations in the future. Getting it wrong, however, can lead to multiyear overruns of large system modernization efforts.

A sound integration architecture could decouple user-facing systems of engagement from the backing core IT systems, thereby reducing dependencies and eliminating sizable tech debt. This often accrues when utilities build business functionality within rather than on top of core IT systems. With this layer of abstraction in place, future M&A efforts could be simplified and consolidations made easier. Applying the architecture principles from the lean governance model at the integration layer could serve as a replicable blueprint for scaling and expanding the platform over time.

We recommend beginning by focusing on use cases within work, asset, or customer management because these core IT systems are central to many digital value cases. Organizations can start with one or two foundational use cases (such as customer payment journey or asset analytics) to prove end-to-end platform integration. For these use cases, the integration patterns can take two forms: operational integrations and analytics. Operational integrations can exist as managed enterprise APIs, designed to provide abstracted, consumable interfaces to read from and write back to core systems themselves. Analytical integrations can serve to build out of the "enterprise data hub," replicating data from core systems—often in streams or real time—for analytical use cases.

Successful delivery of this foundational integration architecture layer requires a combination of strategic guidance from enterprise architecture to help steer teams on key design decisions. Additionally, new or underrepresented skill sets, such as data engineering, may be needed in the organization for it to scale, as well as a sound cloud strategy and infrastructure automation capability.

Phase 2: Establishing consumable interfaces, integration points, and self-service tooling

With the foundational integration patterns in place, the next step comprises the development of consumable interfaces and integration points and the associated self-service tooling rollout. Socialization of the enterprise APIs, data marts, and available integration points through living documentation artifacts (such as Swagger or wiki-hosted data catalogs) could open the doors for business as consuming apps, dashboards, automation bots, and other products begin leveraging the platform's offerings. A self-service model is ideal, where consuming teams have everything they need at their fingertips to find, connect to, and communicate with points of integration across the platform. Sandbox environments (secure, isolated areas for experimentation with data and integrations) could be set up to encourage citizen development—a safe way to explore new use cases for harnessing the data within the platform.

Phase 3: Incrementally scaling to additional domains and expand off-the-shelf accelerators

Once the first end-to-end use cases are delivered for a given domain, phases 1 and 2's processes can be repeated to expand the API and data catalogs with additional business domains, data sets, or system integrations. These could be based on use cases and prioritized by business value. In addition, the integration patterns defined in phase 1 could be applied to additional core systems (such as planning, scheduling, and outage) to bring new read or write capabilities to consuming apps and user-facing products. As the platform delivery initiatives scale and capabilities mature over time, efforts could be focused on the development of assets and capabilities to accelerate platform adoption for consuming use cases. For example, this could include software development kits (SDKs) for easy platform integration or reusable components for engagement layers (including dashboard widgets, mobile and web libraries, forms, and more). Assets like these can accelerate development and help speed up the adoption of tools across business or customer workflows. More advanced acceleration use cases may include cross-platform services such as event hooks or notification services.

Within the data and analytics space, an open-source library of baseline analytics models could help kick-start new teams or inspire new citizen development experiments to unlock untapped value from existing data sources.

Beyond IT: What also has to be true to transform the industry

Building a comprehensive digital utility platform is much easier said than done. Taking key lessons learned from other sectors, we have uncovered six significant factors that could lead to success.

1. *A strong CEO and executive team backed by a board willing to stay the course.* While achieving a better product and service is likely a technical certainty, the path can be rocky. Data privacy, cybersecurity, model drift leading to adverse outcomes, critical talent leaving, and many other issues can derail short-term efforts. Yet an organization with a strong CEO and high-performing executive team committed to the vision will likely overcome such obstacles. Technology talent, however, is vital for a successful team—utilities are often run by engineers, lawyers, and accountants without the necessary technical expertise to guide change.

2. *Product and platform agile operating models combined with lean management principles.* More legacy companies, from automotive to energy, are realizing that they need to adopt a new agile operating model and product development culture. For a successful utility platform to be built, a fundamental change needs to happen across the enterprise, from the front line and back office to executives and the boardroom.

   IT and business silos need to be completely broken down into sustained, impact-oriented product teams, with platforms carefully separated into systems of record versus product that represent systems of insight and engagement. Correspondingly, the product manager's or product owner's role will become more important. A large utility will likely need to hire or train more than 30 product managers—a quick search online across major utilities indicates the current dearth of product manager or product owner roles.

   Beyond digital products and platforms, the rest of the utility organization needs to accelerate more than 30 years of lean management system into three to five years. Here's the opportunity for utilities to move waste and variability, improve frontline problem solving and accountability, enable performance dialogues, and operate a utility with a cohesive operating system—all enabled by lean management system thinking. The combination of lean plus digital is critical: research shows that the hardest part of digital transformations is not talent, technology, or data (although those are difficult enough) but driving operating-model changes that ultimately ensure that the intended business outcomes are achieved.

   To enable this new operating model, winners will hire more employees such as agile coaches, ML engineers, or full-stack developers—people who are in great demand globally. Utilities need to attract this talent by creating compelling career paths linked to the opportunity to build an industry-leading digital platform that drives industry consolidation and plays a meaningful role in the energy transition. Companies could consider "acqui-hiring" a lot of talent at once by buying one or more small software start-ups.

3. *Key differentiators that are built, not bought.* Building a comprehensive digital utility platform is not just about upgrading to the latest management system. Research shows that reliability, safety, resilience, affordability, and customer experience, among others, have to be internally developed for organizations to achieve best-in-class levels of insight and action and industry-leading differentiation.

4. *Domain-based and design-led customer service and workflows.* Domains need to be at the customer care level, electric-distribution asset management, workforce management, and supply chain. Core utility customer journeys (such as paying bills or reporting an outage) and utility workflows (for instance, vegetation or asset management) can be reimagined by using design thinking to create better products and service. While all domains are important, the most critical is getting supply chain right—this will enable industry consolidation. Utilities spend substantial amounts of capital on supply chain but, due to the incredibly fragmented

industry, wield almost no buying power. The winning industry consolidator will most likely have a meaningfully better, digitally enabled supply chain.

5. *Recognizing the importance of cloud.* A 100 percent cloud-based platform that recognizes the flexibility and AI-powered capability of cloud far outweighs any capital-expenditure or operating-expenditure accounting treatment. "The digital utility platform needs to be in the cloud" is not a technical statement anymore. Whether the IT infrastructure is in a data center or the cloud is irrelevant—what is important is that cloud capabilities far surpass what is generally available in an on-premises data center—as shown by the pace of available data science and ML and AI capabilities that major cloud providers (such as AWS, Azure, and Google Cloud) have released in the past few years. Regarding the global utility industry in this space, an important milestone was reached in 2019 when Enel (Italy's national entity for electricity) became the first large utility to be 100 percent cloud-based.[2]

6. *Developing stakeholder skills to enlighten and empower regulators, legislators, the workforce, and other stakeholders.* Historical or legacy regulatory requirements that were mostly put in place in reaction to, or in anticipation of, historical events can hinder progress. Radical transparency is required to educate all stakeholders, partially enabled by a much more rigorous approach to data.

When a utility is transformed by both digital and lean management using the principles above, three things will likely be true. First, the utility could be higher performing (for example, across reliability, resilience, customer satisfaction, safety, and compliance) and more affordable. Second, with a rich focus on data and analytics, the utility could have the ability to prove better performance and cost outcomes to third parties, which would enable an M&A strategy. Third, the digital platform and operating model will be extendable so that acquired utilities could be "bolted on" to improve the utility's performance.

The path ahead

It's not clear yet whether any utility in the North American industry has transformed enough digitally to impact the fundamental industry structure. A number of utilities have digital transformation strategies underway in various forms—some are standing up digital units (for example, using the digital-factory concept), some are systematically upgrading their core systems, and others have made great progress in specific journeys or domains. While it's difficult to estimate exactly how much investment is required, it is likely to be between $500 million and $1 billion. A winner with across-the-board, industry-leading performance and a compelling M&A platform could build a company valued at more than $200 billion—one that consistently delivers top-quartile reliability and customer experience while keeping customer rates among the lowest in the industry.

This is a once-in-a-generation moment. Significant step change in performance could be achieved in a consolidated, regulated industry that is asset intensive, engineering focused, and safety conscious. The first few leaders who recognize the end-to-end opportunity in a greenfield digital utility platform—one that drives an M&A strategy to bring a better and more affordable energy product to millions of customers—could accrue disproportionate value. Why wait?

[2] "Enel 'full cloud': All the advantages of being the pioneers," Enel, July 11, 2019.

Adrian Booth is a senior partner in McKinsey's Bay Area office, where **Aditya Pande** is a partner; **Eelco de Jong** is a partner in the Charlotte office; and **Ben Elder** is a senior director of engineering in the Atlanta office.

A new approach to advanced analytics in utility asset management

Studying how one North American transmission and distribution utility's implementation of advanced analytics in asset management can help other organizations embark on similar journeys.

This article is a collaborative effort by Anjan Asthana, Pranjal Dubey, Alfonso Encinas, Anand Mohanrangan, Aditya Pande, Luis Fernando Rios Siliceo, Jesús Rodríguez González, and Willem van Schalkwyk, representing views from McKinsey's Electric Power & Natural Gas Practice.

Asset management can account for a significant percentage of a transmission and distribution (T&D) company's operating expenses and capital expenditures, with optimized operations and investments key to generating savings. New technologies can enable companies to capture these efficiencies. In fact, a recent McKinsey article explained how T&D utilities can leverage advanced analytics in their asset management strategies to unlock 10 to 20 percent in savings while improving overall reliability and performance of their networks.[1]

This article builds on that thinking and takes a close look at a North American T&D utility, which we refer to as UtilityCo. In 2021, UtilityCo leveraged advanced analytics in asset management to unlock savings of 20 to 25 percent in operating expenses and 40 to 60 percent in capital expenditures, which could then flow as savings into the profit-and-loss (P&L) statement or be reinvested to deliver significant reliability improvement. These savings and increased investment capacity are particularly relevant given today's increasing constraints, including pressure from customers on affordability, inflation growth, supply chain bottlenecks, and the growing need for investments in the energy transition, such as renewable-energy solutions, electric-vehicle charging infrastructure, and cybersecurity. Based on the success of the initial model, UtilityCo developed a road map that scales the asset management risk–based approach to two-thirds of the capital portfolio over two years.

The following case study highlights the results of implementing advanced analytics at UtilityCo, including the approach taken, the lessons learned, and the best practices to adopt for others embarking on a similar journey. Although this article is presented as a stand-alone example, our experience shows that the results from applying advanced analytics to asset management are accelerated when deployed as part of a broader organizational transformation.

UtilityCo: An overview

UtilityCo faced a number of key challenges that are common in the industry. For example, the utility didn't take a risk-based approach when making asset replacement decisions or prioritizing preventive-maintenance (PM) activities, and it had decentralized asset management operations, with each operating company taking a distinct approach and methodology. In addition, although UtilityCo was able to collect valuable data, the data were underused and stored in multiple systems. Finally, UtilityCo relied on rules that oversimplified asset management decisions—for example, the "three strikes" rule, which called for replacing cables after they experienced three outages.

The results from advanced analytics

UtilityCo was able to effectively use advanced analytics in asset management in four ways. First, it optimized capital expenditures either by maintaining current risk and spending less—and letting the excess capital expenditures flow into the P&L or be reinvested to deliver more reliability—or by spending the same amount and achieving higher reliability through replacing the riskiest assets. Second, it lowered PM operating expenses by optimizing PM activities. When successful, this optimization can deliver similar or better reliability at lower cost. Third, it lowered corrective-maintenance (CM) operating expenses by lowering spending on CM after those riskiest assets had been replaced. And fourth, it replaced the riskiest assets to help achieve higher reliability (measured as lower SAIDI and SAIFI[2] performance) due to fewer failures.

Regarding capital expenditures for UtilityCo's transmission transformers, the company underwent a paradigm shift, collecting data about each dollar's impact on interrupted customer minutes. With this new perspective, UtilityCo determined it could reduce risk approximately two to three times over while spending the same amount, maintaining the same capital expenditures, and reducing customer interruptions. Alternatively, UtilityCo could maintain

[1] Rui de Sousa, David González Fernández, Jesús Rodríguez González, and Humayun Tai, "Harnessing the power of advanced analytics in transmission and distribution asset management," McKinsey, April 9, 2018.
[2] System average interruption duration index and system average interruption frequency index.

the same level of risk as determined by the current plan while spending 40 to 60 percent less, thus creating capital headroom for reinvestment, maintaining the same level of customer interruptions but reducing capital expenditures (Exhibit 1). Another option was to select a pathway that both reduced risk and required less spending.

On operating expenses, UtilityCo had the option of spending the same amount on PM—removing 1.5 to 2.0 times the level of risk from the system compared with the current baseline—or maintaining the same level of risk as determined by the current plan while spending 20 to 25 percent less on PM.

For the underground-cables asset class, UtilityCo was able to avoid up to 70 percent more outages as compared with the baseline by replacing its riskiest cables. The optimization model also gave UtilityCo the flexibility to achieve either higher reliability at current spending levels or P&L savings at current reliability.

Finally, UtilityCo developed a visualization platform that displayed the results produced from advanced analytics (see sidebar "Unlocking value as part of a broader transformation"). This dashboard allowed UtilityCo to visualize, prioritize, and implement new maintenance activities (Exhibit 2).

Exhibit 1

The savings potential for UtilityCo's transmission transformers improved by 40 to 60 percent for capital expenditures and by 20 to 25 percent for operating expenses.

Baseline and optimized capital expenditure spend, $ million

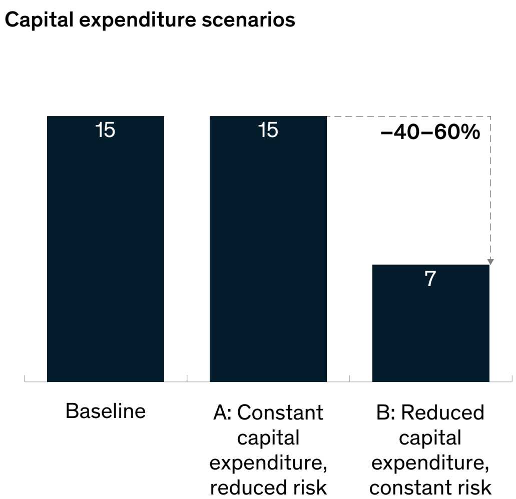

Options for unlocking value:
A: Reduce risk[1] by 2–3× compared to baseline while keeping capital expenditure spend constant
B: Alternatively, reduce capital expenditure by 40–60% while maintaining same risk reduction as in baseline

Baseline and optimized operating expenditure preventative maintenance (PM) spend, $ million

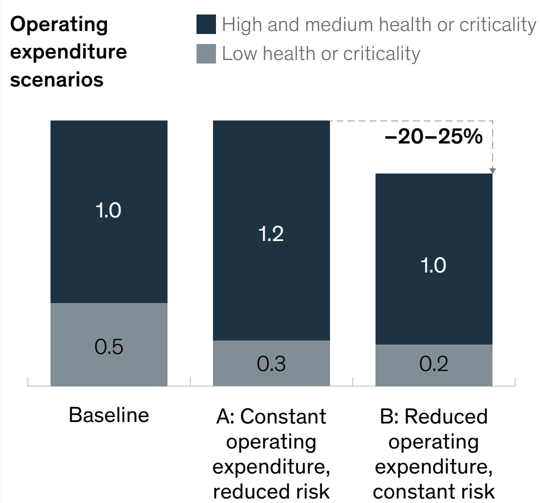

Options for unlocking value:
A: Remove 1.5–2× more risk[1] compared to operating expenditure baseline while keeping operating expenditure spend constant
B: Alternatively, reduce operating expenditure by 20–25% while maintaining same risk reduction as in baseline

[1] Risk is a calculation of health score times criticality.
Source: McKinsey UtilityX analysis

The approach

UtilityCo's advanced analytics–led approach entailed leveraging both internal and external asset data to calculate a health score (the probability of failure) and criticality (the cost of failure) of a given asset. From there, it used the health score and criticality to estimate asset risk and prioritize asset replacement and maintenance activities based on risk (Exhibit 3).

Early on, UtilityCo had a clear plan to scale advanced analytics across all its assets and operating companies. It prioritized assets based on impact, including operating expenses and capital expenditures; time to impact; and feasibility, such as quantity and quality of data and the technical difficulty of building models. In addition, when sequencing distribution assets, UtilityCo followed the concept of "lead" versus "follower" (Exhibit 4).

Exhibit 2

A visualization platform allowed UtilityCo to visualize, prioritize, and implement new maintenance activities.

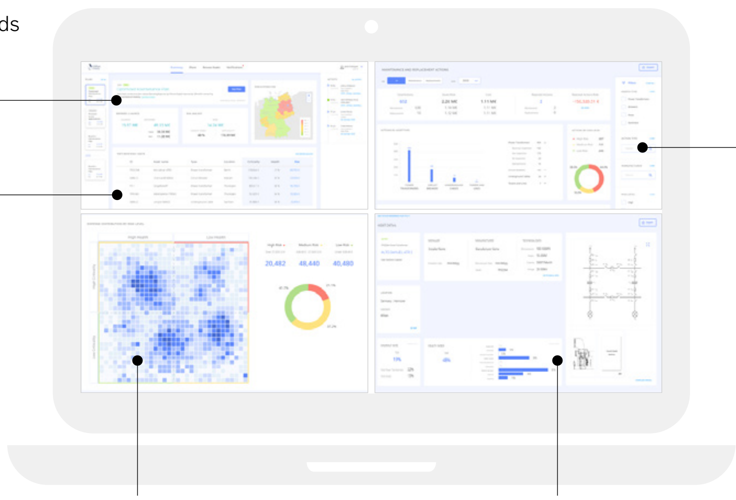

Unlocking value as part of a broader transformation

UtilityCo implemented the asset management advanced-analytics solution as part of a broader business transformation. While the broader transformation wasn't essential to implementing the advanced-analytics solution, it accelerated the adoption process. First, UtilityCo had already adopted a change mindset, so the implementation of a new asset management strategy was more manageable. Second, cross-functional teams were already working on various initiatives—for example, the IT and transmission and distribution (T&D) teams were focused on workforce management—and were thus able to build on these relationships to drive the advanced-analytics initiative. Finally, to successfully scale asset management to other assets and operating companies, UtilityCo needed to fill new technical roles, such as data scientists and engineers, and upgrade its IT infrastructure via cloud migration and machine-learning platforms, among other options. This was easier to do as part of the transformation because UtilityCo was already in the process of building a digital center of excellence, which was able to manage these new functions.

For example, it ensured that poles (lead) were modeled before crossarms (follower).

Health score: Estimating the probability of failure

Depending on the asset, UtilityCo considered more than 100 variables to estimate the probability of failure. A machine-learning model was trained on internal data (such as the age of the asset, work orders, and failure history) as well as on external data (such as weather data, which stretched back a few years). A holdout data set[3] was used to test the model performance. For example, when looking at the transmission transformers asset class, the model was able to predict approximately 45 percent of failures in approximately 20 percent of the data (Exhibit 5).

Exhibit 3

UtilityCo's approach for leveraging advanced analytics was based on the probability and cost of failure as well as on building an optimization engine.

	Description	Outputs
1 **Compute probability of failure (health)**	Calculated the probability of failure for each asset based on the following: • historic outage data • structural data (eg, age and manufacturer) • environmental data (eg, soil temperature and water volume)	Probability of failure for each asset in the following year
2 **Define impact of failure (criticality)**	Calculated the cost of failure for each asset based on the following: • repair • service • safety • environmental impact	Cost of failure for the asset for the following year
3 **High-level clustering for the asset**	Grouped clusters of assets to optimize operational efficiency based on business rules: • eg, grouped assets for replacement at the protective-device level	Clusters of assets
4 **Optimize replacements and inspection**	Generated multiple scenarios to optimize the schedule for the inspection or replacement of assets based on certain constraints; eg, geography and age	Optimized list of assets to be replaced over the next year and prioritized project-management activities

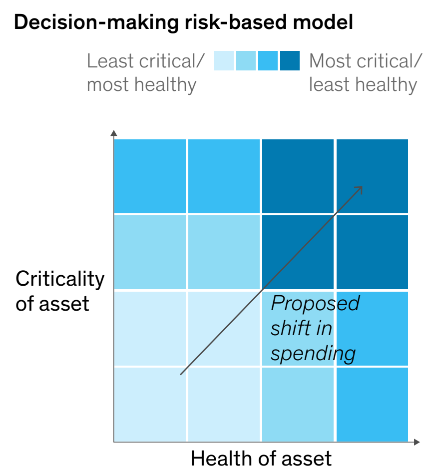

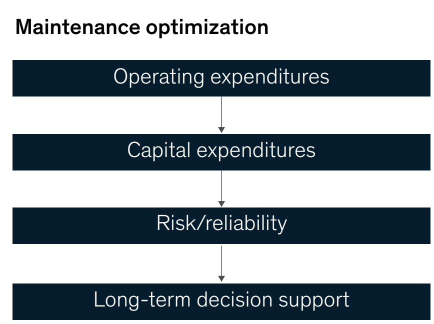

Exhibit 4

UtilityCo's sequence for distribution assets was based on impact and feasibility.

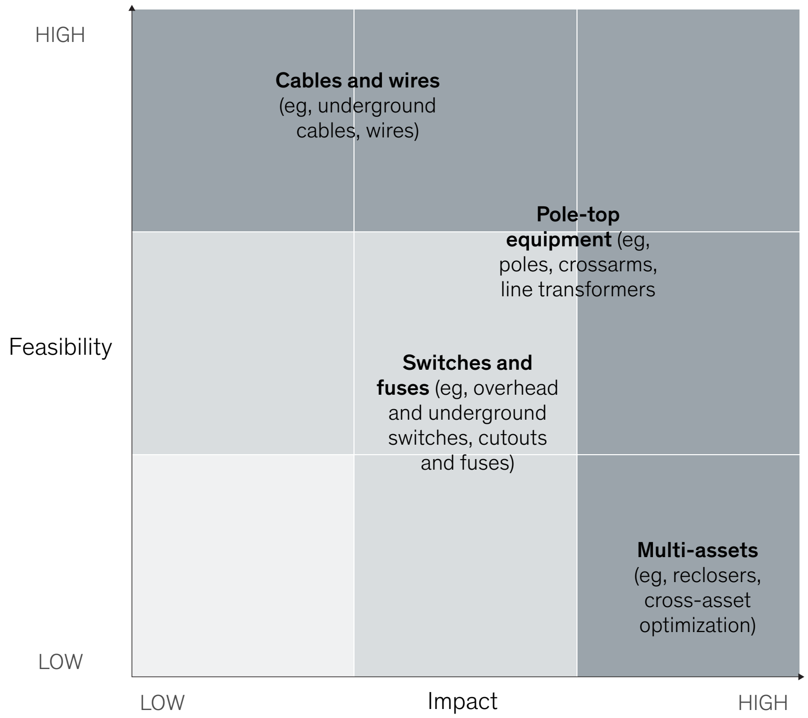

To build the health model, UtilityCo aggregated data from several different systems, such as geographic information and outage management systems. The utility then cleaned and unified the data in preparation for the machine-learning model and identified prediction targets. In some cases, the process was straightforward, such as labeling a transformer that had suffered an outage. In other cases, it was more difficult, such as when labeling a failed cable that was missing a serial or part number. Next, the data were divided into a training set, which was used to train the machine-learning model, and a test set, which was used to help test the performance of the model after training (for example, testing how often the model correctly predicted asset failure). Because UtilityCo was interested in going beyond a standard machine-learning algorithm, it also incorporated a failure-mode analysis into the transmission-transformer models to support detailed assessments of probability of failure by component and to help with the prioritization of condition-based maintenance.

Finally, UtilityCo combined the outcomes from a previous engineering health model with the machine-learning model to calculate the probability of failure and improve performance.

Criticality: Estimating the cost of failure
UtilityCo estimated the cost of failure across several dimensions, including repair, service, safety, and the environment (Exhibit 6). Repair costs are those related to bringing the asset back online after a failure, service costs are estimated based on lost revenue and other factors related to the "importance" of the customer (for example, a hospital

[3] Holdout data, or test data, are historical, labeled data used for validating machine-learning models.

Exhibit 5

Estimating the probability of failure for transmission transformers relies on key input data, important variables, and sample output.

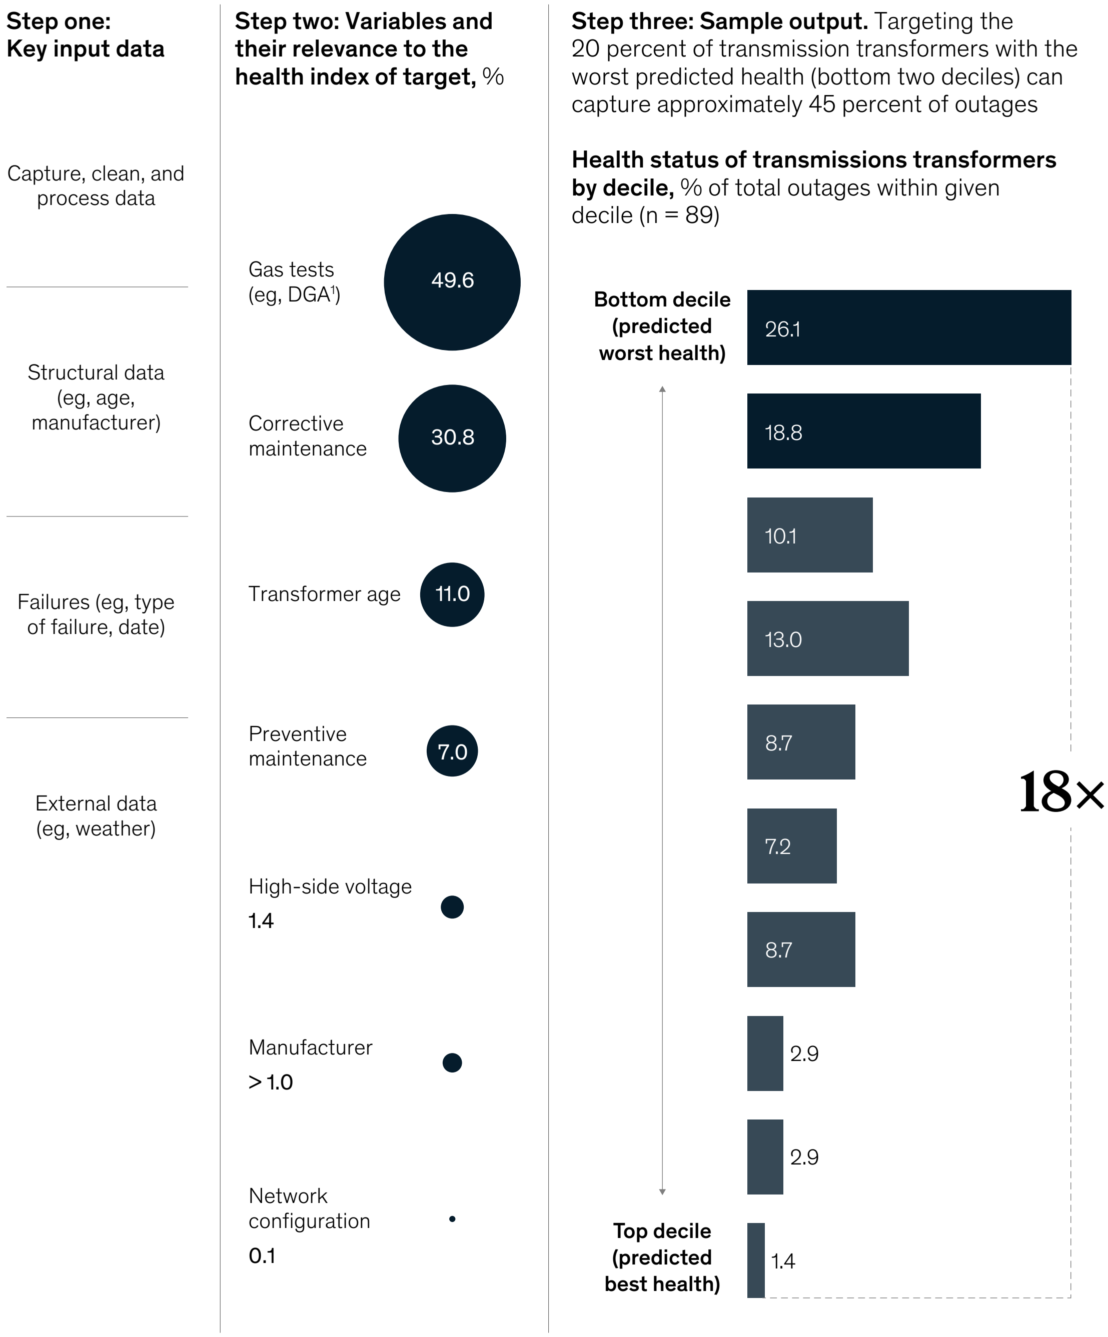

[1]Dissolved gas analysis.
Source: McKinsey UtilityX analysis

is considered "more important" than a single household), and safety and environmental costs are dependent on location and asset type.

Depending on the asset, utilities will need to group the replacement of assets for operational efficiency. For example, UtilityCo found that grouping the replacement of underground cables and applying operational constraints made sense, but it wasn't necessary to cluster the replacement for transformers. Typically, simple business rules—such as clustering cables to be repaired based on the protective device they are connected to—can be used to cluster assets.

Building an optimization engine
UtilityCo used an optimization engine to prioritize asset replacement and PM activities based on the risk of the asset. To build the optimization engine, UtilityCo first estimated the risk of each asset by multiplying its health score by its criticality. Asset replacement was then prioritized based on the resulting risk score and the asset's replacement cost. For example, a high-risk transformer with a lower replacement cost was prioritized over the same type of transformer with a higher cost. To optimize PM, a detailed failure-mode analysis was incorporated into the optimization engine to enable estimates of how much risk was removed by each PM activity. In addition, the optimization engine factored in the cost required to perform each of these activities and prioritized activities that reduced the most risk at given costs.

Exhibit 6

Key dimensions are used in calculating the cost of failure for a typical asset.

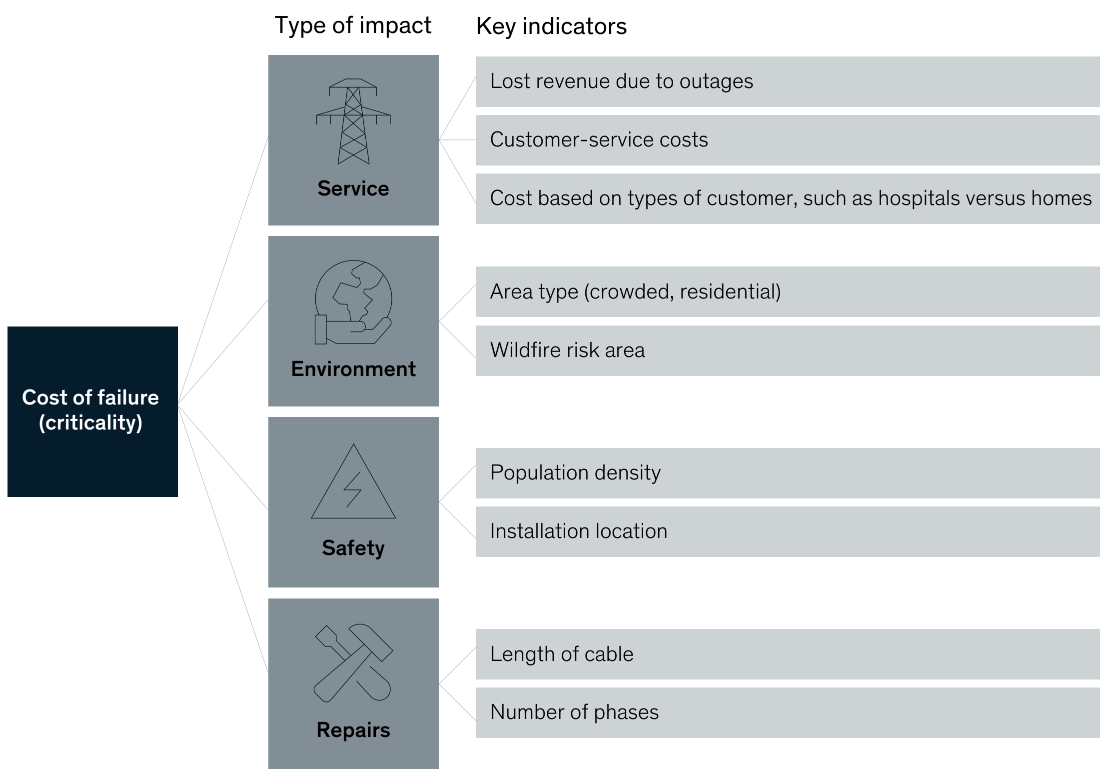

To optimize PM, a detailed failure-mode analysis was incorporated into the optimization engine to enable estimates of how much risk was removed by each PM activity.

Lessons learned and the recipe for success

Implementing an advanced analytics–led asset management program resulted in several lessons learned, particularly with regard to getting started, data quality, talent and capabilities, change management, and implementation and governance.

First, the team faced internal resistance when getting started, including concerns about not having enough data, or the right data, to address regulatory considerations. The first key step to addressing this resistance was to implement a proof of concept, identifying assets that had good enough data to get started and developing a solution that was better than the current state. Success with the proof of concept gave UtilityCo the confidence to proceed with rolling out the solution across multiple assets and operating companies.

Another lesson involved data that were either siloed, scattered across several different systems, or incomplete or duplicated. For example, data from one asset class were missing installation dates. As a workaround, the manufacturing date was used instead. Developing data architecture and putting processes in place to capture and perform quality control checks on the right types of data were key to addressing this issue going forward.

Although the artificial-intelligence and machine-learning space is still emerging, it is growing quickly. The data scientists and engineers who are key to building solutions are scarce. Thus, UtilityCo built a digital center of excellence and leveraged it to manage the pipeline of talent and to develop processes and trainings to drive consistency across the organization.

During the advanced-analytics implementation, UtilityCo asset managers were asked to make changes to their management processes (see sidebar "Driving additional value from cross-asset optimization"). The key to addressing this issue was to engage the asset managers early on and bring them along as the solution was being developed.

Finally, incorporating advanced analytics into the processes for selecting assets for replacement and updating maintenance processes and policies based on model recommendation was an implementation and governance challenge. Engaging subject matter experts (SMEs) early and running pilots to test new processes gave UtilityCo confidence in its models' abilities to meet its needs.

Several key ingredients contributed to the success of UtilityCo's asset management transformation:

— *Top-down leadership buy-in and push.* Because implementing such a solution usually involves changes in processes across different departments, entailing new technologies, UtilityCo discovered that leadership buy-in and

Driving additional value from cross-asset optimization

One of the key advantages of asset analytics is cross-asset optimization. Once UtilityCo implemented advanced analytics for multiple assets, it was able to take a risk-based approach that allowed replacement capital expenditures to be invested in the assets that reduced risk and increased reliability the most.

Because risk measurement is the same across all assets, the risk-based approach makes it easier to decide where to invest the next million dollars of capital expenditures to achieve company-wide objectives. For example, with the objective of improving reliability, UtilityCo determined it was better off spending more on one asset over another based on estimated system average interruption duration index (SAIDI) improvement per million dollars spent, resulting in approximately 94 percent improvement to reliability (exhibit).

Exhibit

Everything being equal, investing replacement capital expenditures in asset 11 instead of asset 1 will lead to an improvement in reliability of approximately 94 percent.

Estimated SAIDI[1] improvement, $ million/SAIDI-minute improved[2]

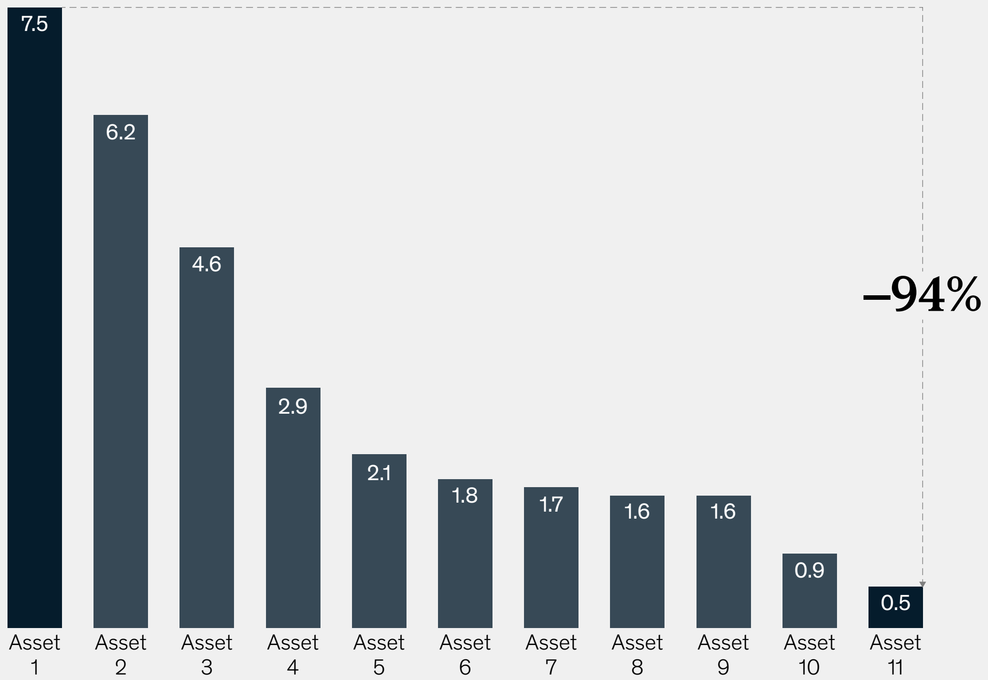

Note: Data sanitized by taking 90 percent of the original value.
[1]System average interruption duration index.
[2]Numbers have been rounded.
Source: McKinsey UtilityX analysis

push from the top was critical for advanced-analytics adoption. For example, the vice presidents of T&D and IT joined biweekly sprint reviews to take stock of progress, encourage the team, and communicate the importance of the effort.

— *Agile approach to working.* By setting one- to two-week sprints and working collaboratively to achieve the goals for each sprint, the team had time to cocreate, bring along team members, and course-correct as needed without falling behind. As part of this, a cross-functional team of business and IT was key to making the implementation successful. For example, in a joint meeting with IT and business, an SME was able to provide guidance on which corrective work orders should be included in the data used to calculate the probability of failure for a given transmission transformer.

— *Not letting perfection be the enemy of getting started.* It was more important for UtilityCo to get started with the data that were available and demonstrate some economic impact than to get everything perfect (or build a big data lake) before starting. For example, UtilityCo started with proof of concept for two assets to demonstrate the value before scaling to more assets.

— *Change management.* Asset managers typically have years of experience managing assets and follow specific procedures and policies. As a result, UtilityCo found it difficult to convince its managers to take recommendations from machine-learning models. For example, when it came to replacing underground cables, UtilityCo asset managers were accustomed to the three-strike rule—in which a cable is replaced after three failures within a 24-month period—and it was a difficult task to convince them to replace a cable that had a high probability of failure but no previous failures. One successful strategy was to involve the team from the beginning and bring them along in the process of building the models.

UtilityCo transformed its asset management strategy from a manual process that used very limited data to make asset replacement decisions and a preventive, one-size-fits-all maintenance approach into a strategy that leverages extensive data and advanced analytics to make replacement and repair decisions based on asset risk. In making this change, UtilityCo unlocked significant savings in capital expenditures and operating expenses while increasing reliability for customers and regulators.

Anjan Asthana is a senior partner in McKinsey's Miami office; **Pranjal Dubey** is a consultant in the Bay Area office, where **Anand Mohanrangan** and **Aditya Pande** are partners and **Willem van Schalkwyk** is an associate partner; **Alfonso Encinas** is a partner in the Washington, DC, office; **Luis Fernando Rios Siliceo** is a consultant in the Monterrey office; and **Jesús Rodríguez González** is a senior partner in the Madrid office.

Smart scheduling for utilities: A fast solution for today's priorities

AI-driven schedule optimizers are alleviating long-standing headaches for utility companies by reducing employee downtime, improving productivity, and minimizing schedule-related service disruptions.

by Sohrab Rahimi, Zachary Surak, Jackie Valentine, and Akshar Wunnava

Utilities today are squeezed across multiple priorities—including reliability, cost, and safety—and are facing increasing challenges related to labor shortages, regulatory scrutiny, and a post-COVID-19 hybrid work environment. Companies have generally taken on efforts to address their priorities individually, and this one-metric focus often leads to the inaccurate perception that there is a trade-off among reliability, cost, and safety. As a result, significant optimization improvements remain for those who are ready to take a more holistic, end-to-end approach.

Our experience in process transformation efforts across utilities indicates that the largest drivers of execution waste relate to the initiation of work orders, planning and scheduling handoffs, and information silos. Many of these issues can be traced back to traditional work management processes, which rely heavily on many time-consuming and inconsistent manual processes.

Smart scheduling involves analytics-powered algorithms and user-centric interfaces that can be deployed in a matter of months and within existing systems to build better, faster schedules. AI-enabled smart scheduling that efficiently matches resources with work can transform companies' ability to drive long-needed improvements across multiple competing priorities. It can free up scheduler time, boost worker utilization, and increase productivity by 20 to 30 percent. These additional resources can then be used to reduce overtime, insource contractor spend, or reduce job backlogs.

In our experience, successful deployment of smart-scheduling tools requires utility companies to learn five key lessons:

— Data are crucial but should not be a barrier to starting.

— Technology must work in conjunction with processes.

— Businesses must clearly specify their optimization criteria.

— Piloting, followed by intentionally scaling, a "light tech" scheduling solution is vital to increasing adoption.

— Solutions must be user-friendly and holistic.

Deploying new technologies can significantly improve scheduling

In a previous, industry-agnostic article, we laid out how optimizing work management—starting with smart scheduling or scheduling optimization—can improve grid reliability, the efficiency of capital deployment, cost, safety, and employee engagement.[1]

AI-enabled smart scheduling can transform companies' ability to drive long-needed improvements across multiple competing priorities.

[1] Jorge Amar, Sohrab Rahimi, Nicolai von Bismarck, and Akshar Wunnava, "Smart scheduling: How to solve workforce-planning challenges with AI," McKinsey, November 1, 2022.

Schedule optimization, however, is one of the most challenging optimization problems due to variations in types of work and operations. This variation makes solutions hard to generalize and therefore hard to scale. Additionally, the mathematical complexity of optimization equations and the number of decision variables mean models take a long time to run. To be truly useful, optimization models need to operate in almost real time so that they can react to changes such as employee sick days and unexpected demand surges.

While classic optimization models have been around for decades, the advent of new technologies in AI and cloud infrastructure allows for the rapid development and deployment of tools that bring deep analytics and optimization engines to the scheduling process. These tools have also reduced the cost of deploying an end-to-end schedule optimization solution and can sit on top of existing work management systems. Additionally, using AI improves the quality and functionality of scheduling in a number of ways:

— offering the most optimal solution given a range of interdependent constraints and dynamic, ever-changing demand

— providing a consistent, systematic approach with no human bias

— delivering significantly faster computation than manual processes, which improves the ability to adapt to unexpected changes in operations

— lowering HR requirements, which frees up capacity to focus on other areas

Smart scheduling offers benefits for utility companies

For electric and gas utilities, scheduling is a central function that matches demand for services with the crews, materials, and equipment needed to perform those services. Utilities have a variety of different work types—including emergency jobs, short-cycle jobs, and long-cycle jobs—with varying scheduling dynamics. Smart scheduling provides benefits for each work type:

— Emergency jobs have high importance but low predictability and may require a crew to be immediately reallocated from another work site. These schedule "break-ins" require real-time juggling of crews and often cause churn and rework for schedulers. Smart scheduling can help block off capacity for these emergent break-ins via dynamic schedule loading. For example, only 60 to 70 percent of capacity may be allocated in a given week if algorithms predict, based on historical data, that 30 to 40 percent of time will need to be spent on emergency jobs. Smart scheduling can also identify the optimal crew to address the emergency job based on factors such as geographic proximity and the priority and state of the crew's current job.

— Short-cycle jobs can typically be completed within the day. They range in complexity: some jobs may require one crew for an hour or two, while others—such as hydro-vacuum excavation—may require several crews for a full day alongside coordination with third-party contractors. The scheduled duration for a short-cycle job may often be several hours more or less than the actual requirement, leading to either a schedule backlog or underutilization. Smart scheduling can better estimate the durations of these jobs using a combination of historical performance and factor-driven adjustments. For example, data on local soil composition can be used to estimate the time needed to dig.

— Long-cycle jobs may require multiple days to complete, and the main challenge is to ensure continuity by scheduling the same crews for the whole duration. These jobs often come with multiple crews and pieces of equipment, plus third-party contractors, which means that smart scheduling can ease the significant mental burden on schedulers.

Schedulers need to coordinate the availabilities of crew, materials, and equipment ahead of time to ensure that all components are ready on the day when the work is to be done. Depending on the type of job, schedulers may need to create crews of different sizes—generally one to four full-time equivalents. Additionally, crews may be qualified only for certain types of jobs, and some jobs (particularly those related to electrics) may also require materials that are not in stock and that have a long lead time once ordered. Most gas jobs, on the other hand, can be done with the materials readily found on trucks. Finally, jobs may require special equipment such as backhoes or diggers.

One of the largest pain points for crews is job delays or "false truck rolls," which occurs when a job cannot be started or completed on time due to the unavailability of the right crew, materials, or equipment. Smart scheduling can help ensure all job components are ready before jobs are incorporated into the schedule.

The tangible benefits of smart scheduling for a US utility

In our previous article, we laid out the significant, tangible benefits accrued by a US electric and gas utility after it piloted a machine learning–based schedule optimizer[2]:

— *Lowered HR requirements for scheduling.* Scheduler productivity increased by 10 to 20 percent, which is the equivalent of freeing up one to two scheduler hours per day.

— *Increased automation for flexibility.* AI models automate initial schedule builds and ongoing optimization and can react to changes in the system (for example, COVID-19, seasonalities, or workforce changes) within one to two days. Manual schedulers may take much longer to adjust to such shifts.

— *Reduced waste.* Over the six-week pilot, dynamic schedule loading and a decreased number of prematurely scheduled jobs meant that break-ins were down by 75 percent, job delays by 67 percent, and false truck rolls by 80 percent.

— *Increased crew utilization and field productivity.* Prior to the pilot, crew members at one of the utility's sites spent 44 percent of their time actually working on jobs (as opposed to being unassigned, training, or traveling). In the automated, optimized schedules, crews could expect to spend 65 percent of their time on jobs. Overall, the pilot achieved an approximate 20 to 30 percent increase in field productivity (Exhibit 1).

Five lessons for utility players in deploying a smart-scheduling solution

Based on our experience, there are five core lessons to keep in mind during the development and deployment of smart-scheduling solutions in a utility context.

1. Data are crucial but should not be a barrier to starting

Many utilities often delay analytics-based scheduling efforts due to a lack of trust in data quality. Most leaders have a misperception that data need to be rich and easy to digest to begin to get value from AI-based tools, but the opposite is true: a small amount of data can yield disproportionate insights. In fact, new data-processing methods can take existing data and make them usable for AI models. To achieve optimal results, utility players will need to map their data landscape and find resolutions to any issues that compromise the quality or usability of the data (see sidebar, "Mapping the data landscape"). This process frequently highlights the relative importance of specific data that can then be prioritized for better data governance and stewardship, which can further increase the accuracy of AI outputs.

These processes can be conducted in as little as three weeks, but space must be built into any smart-scheduling rollout timetable. This time is used to prepare and process data related to timesheets, HR, and job backlogs, as well as to evaluate data quality and to run preprocessing modules to prepare data sets to be used by the AI engine.

[2] "Smart scheduling," November 1, 2022.

Exhibit 1

Smart scheduling at an electric utility improved field productivity and reduced waste by 20 to 30 percent.

Improved field productivity and reduced field waste over a 6-week pilot

80%
reduction in false truck rolls

75%
reduction in break-ins

67%
reduction in job delays

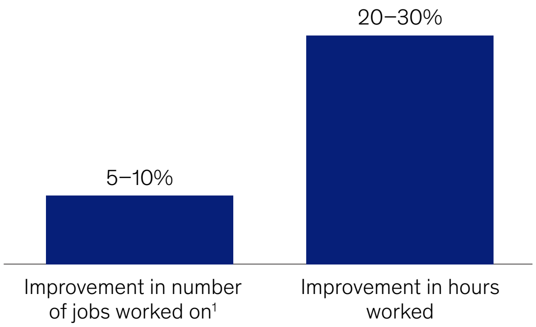

[1]Increase occurred during peak training time, Omicron variant outbreak, and winter weather.

McKinsey & Company

Mapping the data landscape

Utility players will need to map their data landscape—that is, understand data quality and relationships between data sets—to identify resolutions to any potential data issues.

Potential data and operational issues

— Changes to work orders are not tracked over time.

— Schedules are not locked, making tracking adherence challenging.

— Travel time may not be accurately coded.

— Job duration estimates are inaccurate.

— Too many data sources can be overwritten by human inputs.

Potential resolutions

— locking schedules to allow for better metric tracking

— updating job duration estimates

— estimating travel time based on typical patterns

— identifying unknown gaps in crew timesheets to improve quality of crew metrics

2. Technology must work in conjunction with processes
Smart scheduling will be effective only if it works for the end user. Therefore, new technologies must be developed in conjunction with efficient scheduling processes, which should be codesigned with frontline employees. Digital tools, such as smart-scheduling engines, codify the underlying processes, meaning that organizations that do not optimize their processes in tandem with the development of technical tools are at risk of codifying inefficiencies.

Getting the most out of new digital tools may require some or all of the following process improvement initiatives:

— clear job readiness checklists that take into account the specificities of each electric or gas job

— break-in management processes that reduce nonemergency break-ins and quickly reorder the schedule to address emergencies

— efficient handoff meetings to align stakeholders such as schedulers, field supervisors, and warehouse managers

— prejob walkthroughs to ensure site readiness

Additionally, the successful deployment of digital tools requires continuous maintenance of the technical models. This work will require a number of different skills profiles, including capable data scientists, data engineers, and cloud engineers.

3. Businesses must clearly specify their optimization criteria
A smart-scheduling engine can optimize frontline schedules based on several evaluation criteria. For example, the engine could maximize the number of jobs scheduled, minimize operating costs related to shifts or travel time, or maximize service levels by reducing customer wait times. To achieve the best results, it is imperative that business leaders feed clear objectives into their smart-scheduling engine.

4. Piloting, followed by intentionally scaling, a 'light tech' scheduling solution is vital to increasing adoption
Smart-scheduling solutions can be developed as "light tech" overlays on top of existing systems and do not require platform overhauls ("heavy tech"). Algorithms can often be tested in an isolated testing environment to pilot the efficacy of a scheduling optimization solution. This piloting, which should be done in conjunction with schedulers, is essential to train the model for the specific utility company and context. For example, variations in regulations or

To achieve the best results, it is imperative that business leaders feed clear objectives into their smart-scheduling engine.

union-specific requirements can have a significant impact on the details of an optimized schedule.

A key metric during the pilot period is the frequency of manual schedule overrides by schedulers. These overrides can indicate an issue with the underlying model and should therefore happen as seldom as possible. However, some manual intervention will always be required to address last-minute contextual changes such as sick days or employee holidays.

In our experience, it takes at least four to six weeks for smart-scheduling algorithms to reach a 70 to 80 percent match with the final schedules previously created by schedulers, as measured by the percentage of jobs and crew pairings that are the same in each (Exhibit 2). While an optimal schedule is unlikely to exactly match the existing manual schedule, a relatively close match is a good indication that the new algorithm is factoring in the right parameters and will not require frequent manual overrides.

Pilots can also be an important way to build support for the new scheduling methods within the organization, which can make the subsequent rollout easier. In the US utility example used above, schedulers—who were spending four to seven hours a week building and updating the manual schedule—saw that the new technologies could build automated schedules that closely matched their own within minutes.

After a successful pilot, it is important to execute a well-thought-out scale-up plan. This plan should take into account factors such as overall deployment speed, deployment across work types (that is, there may be different considerations for electric versus gas jobs or for short-cycle versus long-cycle jobs), the differing challenges of rural and urban service centers, and resourcing the scale-up (for example, potentially hiring change champions or trainers). Tools and processes can be scaled up in an agile fashion because making iterative improvements over time is generally preferable to trying to perfect the algorithms during the pilot period.

Exhibit 2

Schedule optimizers can improve to build 70 to 80 percent of the final weekly schedules in just four to six weeks.

Typical improvements in scheduling by AI-driven schedule optimizer

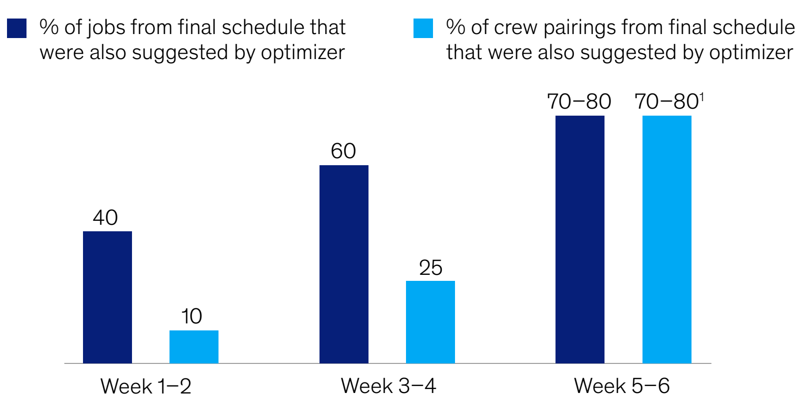

[1]Improvement of about 50 percentage points from week 3–4, driven by favored crew pairings.

McKinsey & Company

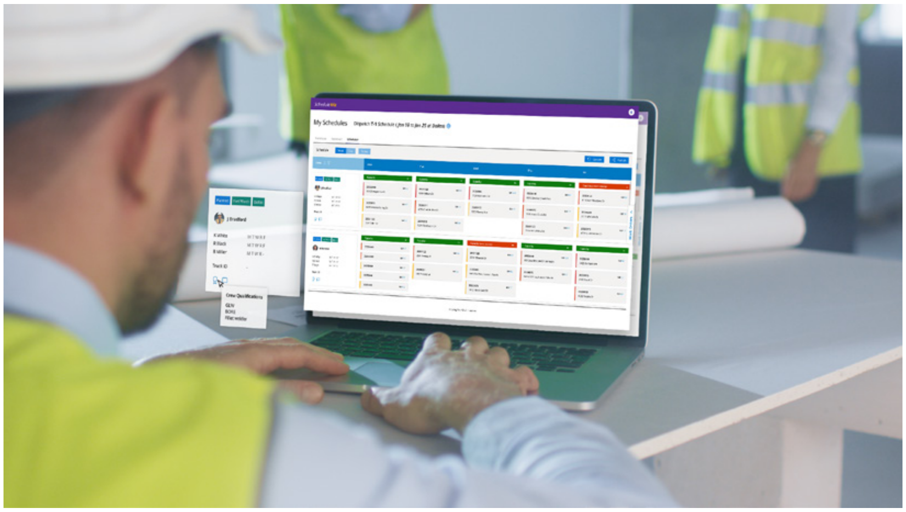

Scheduling interfaces could have several user-friendly features.

5. Solutions must be user-friendly and holistic

To operate successfully, schedule optimization needs to be integrated into a user-friendly, end-to-end digital solution. The final, holistic solution must update constantly, forecast accurately, and incorporate an easy-to-use, interactive front-end interface. Many schedules are currently based in Microsoft Excel, and the benefits of the automated, AI-optimized schedule can be multiplied if companies can incorporate features such as daily or weekly drag-and-drop schedules and a metrics dashboard.

Scheduling interfaces should incorporate user-friendly features, which could include the following:

— preloaded, optimized schedule

— prioritized work orders

— simple and real-time edits

— precise information display

— flexible crew management

Inflation, supply chain issues, and ongoing labor disruptions are making work optimization—which has long been one of the most challenging problems for consumer-facing industries such as utilities—even more complex. When deployed thoughtfully as part of a holistic solution, AI-driven schedule optimizers can significantly improve work management processes, smooth out operations, and boost overall productivity.

Sohrab Rahimi is an associate partner in McKinsey's New Jersey office, where **Zachary Surak** is a senior partner; **Jackie Valentine** is a partner in the Washington, DC, office; and **Akshar Wunnava** is an associate partner in the Dallas office.

The authors wish to thank Jorge Amar, Adrian Booth, Alfonso Encinas, Keith Gilson, and Nicolai von Bismarck for their contributions to this article.

How utilities can use advanced analytics to elevate customer experience

Industry players can both lower costs and improve performance by capturing the full value of today's technologies.

This article is a collaborative effort by Bobby Dean, Arpit Goenka, Vinay Gupta, Nimish Jain, and Humayun Tai, representing views from McKinsey's Electric Power & Natural Gas practice.

Across industries, customers increasingly expect interactions with companies to be simple, predictive, and seamless across channels. Three examples demonstrate this kind of enhanced customer experience:

— A North American utility proactively notifies customers about higher bills during the winter and provides customers with tips on how to reduce their bills.

— A bank alerts its customers about potentially fraudulent transactions and can replace debit cards within 24 hours.

— An airline provides a seamless customer experience when customers move from its app to interactive voice response (IVR) to an agent.

The common thread in these examples is the role that advanced analytics plays in enabling exceptional customer service across channels. These elements can improve customer satisfaction (an increase of 10 to 20 percent) while unlocking lower cost to serve (cost savings of 20 to 30 percent in some cases). For these reasons, companies are seeking to step up their customer service offerings to meet the heightened expectations of customers.

Capturing cost savings and improving satisfaction don't necessarily require complex use cases. Targeted use cases, such as speech analytics, can boost performance by 10 to 20 percent by enabling more effective agent coaching. And IVR analytics can improve containment rates by 5 to 10 percent.

Because many utilities have yet to implement advanced analytics in customer care at an enterprise level, we typically recommend proceeding with a two-phase approach. Utilities should first ensure they have a robust foundation of organizational capabilities and acumen. Then they can explore advanced techniques that enable next-generation use cases to unlock greater levels of cost savings and improved customer satisfaction.

Phase one: Laying the foundation

To begin integrating advanced analytics, utilities first need to identify and prioritize use cases, aggregate the necessary data, establish cross-functional teams, and adapt operations.

Identify and prioritize use cases

Multiple existing advanced-analytics applications have the potential to improve customer service and productivity. But different use cases do not generate equal value, and without guidance, analytics teams typically will not know which use cases to prioritize. To avoid leaving significant value on the table, utilities should ensure their analytics teams work closely with the business to assess a list of use cases based on the following criteria:

Potential value. Estimating the value of a use case should include considerations such as a clear view of the baseline, the value at stake, implementation costs, and the time to impact.

Technical readiness. The organization's technical maturity—encompassing elements such as the availability of data and resources, deployment considerations, the likelihood of adoption by end users, and scalability—can dictate how quickly a use case can generate value.

Strategic fit. A use case should be aligned with the organization's overall goals, particularly regarding its potential to improve customer satisfaction.

Aggregate data for prioritized use cases

Disconnected data sources often lead to inaccurate or incomplete insights. For example, the combination of data from the customer relationship management (CRM) and automated call distribution systems can unlock a host of use cases, including "call propensity by customer segmentation." Utilities can gain a tremendous advantage by harnessing the power of all available data sources for prioritized use cases, thus serving customers more effectively. They must also develop the ability to connect data from different sources using unique customer identifiers.

Utilities should start with readily available data sources to prove the value before adding other data sources to strengthen the use cases. For example, a customer journey use case would require aggregating customer data from multiple channels, identifying the break points, and using these insights to design future interventions. One utility that embarked on this use case augmented contact-center data with data from interactions across the web, mobile, and IVR (Exhibit 1). Analysis revealed more than 50 percent of the customers were not using digital channels prior to reaching out to the contact center, with app usage particularly low. These insights led the utility to focus on increasing customer awareness of its digital channels and revamping some existing digital functionality, such as simplifying bill formats, using visuals to show bill trends, and reducing the number of steps needed for a customer to start a service.

Build a cross-functional team

Many analytics teams work in silos (such as an analytics center of excellence), which means they lack an understanding of the organization's greatest pain points and miss out on receiving regular feedback from the business on the applicability of the insights they generate. For example, to get the most from analytics-driven models that draw on employee data, the analytics team would need feedback from the managers in charge of hiring, recruitment, and operations. Close collaboration among multiple teams is critical.

Exhibit 1

One utility redirected more customers to digital channels with data from interactions across the web, mobile, and interactive voice response.

Illustrative

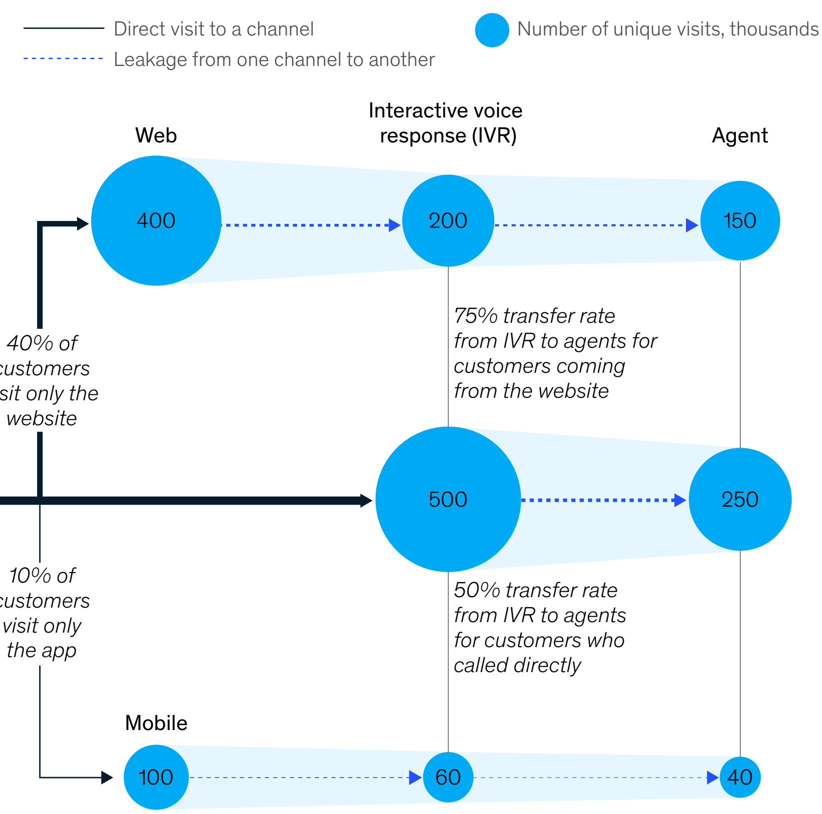

One US utility was able to identify and fix common customer break points in its IVR by promoting a collaborative approach. The analytics team worked with the business to identify customer drop points and closely coordinated with the IT team to make quick adjustments to the IVR. The team then tracked the impact of these changes on customer behavior. The analytics team was able to gauge how frequently customers used self-service features to identify the most- and least-popular nodes and reorder the options accordingly (Exhibit 2). Most were related to problems with technology workflows on the back end or customers being confused by the verbal prompts in certain nodes. By itself, the analytics team could not have identified such adjustments.

Adapt operations

Capturing value from customer-care use cases requires businesses to make appropriate operational changes. For example, increased digital adoption through targeted customer education would decrease the volume of live calls to the contact center, reducing the number of frontline staff required. Operations should coordinate with the workforce management team to ensure the team revises staffing projections to capture the total value of the use case. Analytics

Exhibit 2

Interactive-voice-response analytics reveals specific pain points and opportunities for improvement in the customer journey.

What do payment callers do in interactive voice response (IVR) and before calling?, % of total callers

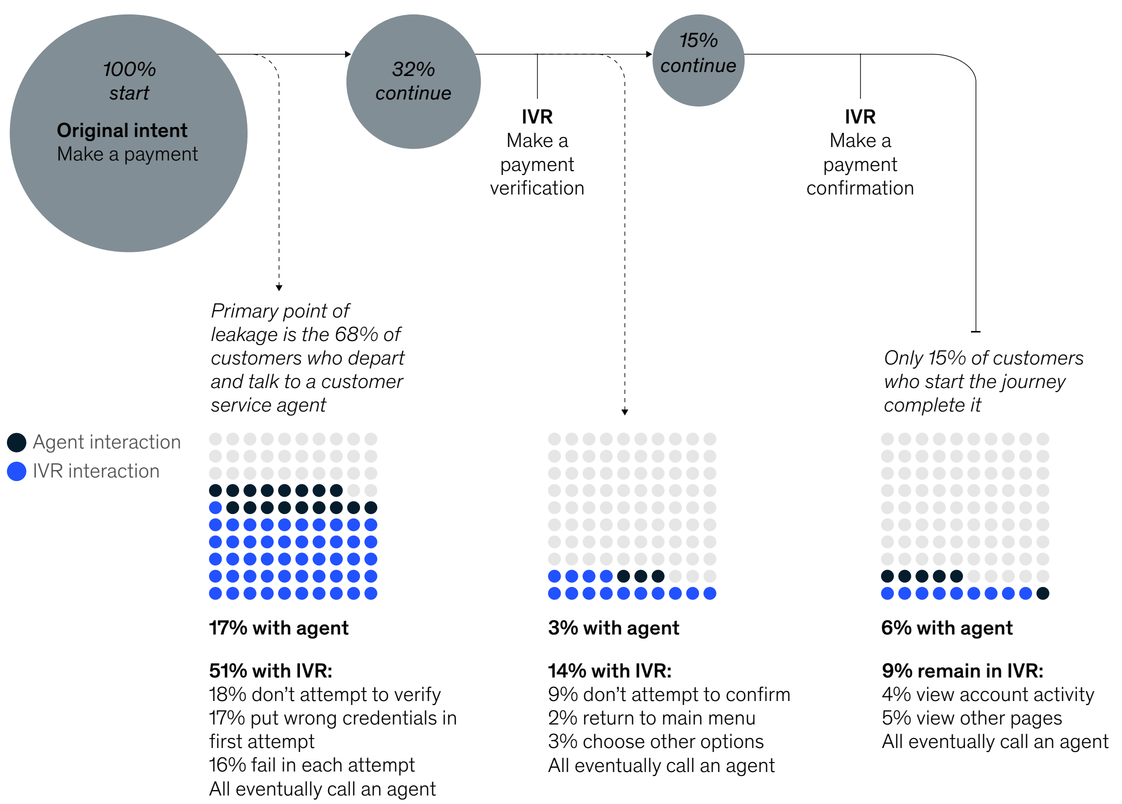

Once utilities have mastered the foundational elements of analytics, they should prioritize the next generation of use cases.

can also play an essential role in tracking KPIs to ensure the use case is generating the anticipated value from initial estimates.

Phase two: Next-generation use cases

Once utilities have mastered the foundational elements of analytics, they should prioritize the next generation of use cases. This process typically requires brainstorming with the business to come up with a list of new use cases, capturing incremental data such as text transcription data from voice recordings, and enhancing their understanding of advanced-analytics techniques. We discuss below two examples of these techniques: machine learning (specifically to build predictive-intent models) and natural-language processing (specifically for speech analytics).

Predictive-intent model

Anticipating why a customer might call is not only a key to excellent service but also the future of customer care. To gain this capability, utilities will need to build algorithms that use advanced machine-learning techniques and data from previous interactions across channels (such as email, chat, apps, and IVR) to forecast call intent.

For example, one utility deployed a predictive-intent model to reduce calls related to unexpectedly high bills. The algorithm estimates the individual customer's upcoming bill and calculates the deviation of the bill from the amount billed the previous month and year. If the deviation is significant, the utility proactively sends a notification to the customer and explains why the upcoming bill could be higher. Similarly, we have also seen predictive-intent models used to adapt IVR to proactively resolve customer concerns.

Today, these models commonly deliver meaningful value across industries outside the utility space. For example, airlines predict call types based on the customer's most recent communication and queries from previous calls; if the purpose of a customer's past few calls involved time of departure, that option is listed first in the IVR. Utilities can learn from successes in other industries.

Speech analytics

The combination of good call recording, advanced speech-to-text conversion, and text analytics (another use of machine-learning algorithms) can generate extensive call-based insights. Specifically, algorithms can pinpoint additional opportunities for call resolution, and speech analytics can support use cases such as agent coaching, call drivers, customer satisfaction drivers, and the automation of quality assurance.

One utility built a speech analytics use case to identify agent-coaching needs by call types. The traditional approach was much more manual, with a handful of supervisors listening in on calls. By harnessing analytics-driven insights, supervisors could understand agent needs at a much more granular level and engage with agents more meaningfully in coaching sessions (Exhibit 3).

Exhibit 3

With speech analytics used to identify agent-coaching needs, supervisors can better understand where their agents can improve.

Example agent performance by quartile for each call type

Agent 1 is a high performer and performs in the top quartiles overall across various categories—but still has opportunities to grow

Agent 1:
Overall quartile 2

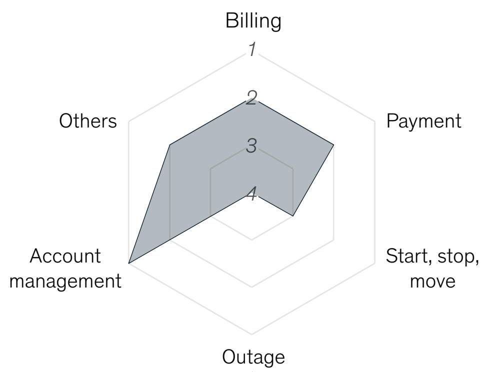

Agent 1 struggles specifically in outage calls; coaching should target this area

Agent 2 is a low performer; and managers should dedicate more time on training this specific individual

Agent 2:
Overall quartile 3

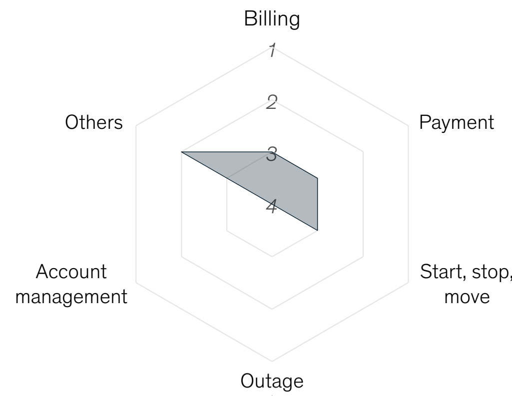

Agent 2 requires the most coaching on outage and account management

Many utilities aren't capturing the full value of advanced analytics today due to a lack of prioritization, limited ability to harness multiple data sources, and limited collaboration across multiple departments. Developing these capabilities will be critical for any utility that expects to be a customer experience leader and provide better service at a lower cost.

Bobby Dean is an associate partner in McKinsey's Washington, DC, office; **Arpit Goenka** is a consultant in the Gurugram office; **Vinay Gupta** is a senior expert in the Waltham office; **Nimish Jain** is a partner in the Southern California office; and **Humayun Tai** is a senior partner in the New York office.

The authors wish to thank Rohit Agarwal and Nitish Gupta for their contributions to this article.

Three ways energy providers can boost resilience and digital customer experience

Electric power and natural gas organizations are at an inflection point, facing a convergence of challenging global trends in market and customer dynamics. Here's how they can boost resilience.

by Bobby Dean, Oliver Ehrlich, Hongqiao Lu, and Scott Perl

Utilities and energy retailers face strong headwinds globally. Tough market dynamics are increasing pressure on customer wallets, driven by rising inflation, energy cost volatility, and energy supply-chain disruptions. In addition, customer needs and preferences around decarbonization—such as clean energy and electrification—continue to grow, forcing utilities and energy retailers to evolve their portfolios to meet these needs. Supply alternatives, affected by declining technology costs for distributed solar and battery storage, among others, have pushed customers to reassess their energy choices at an accelerated pace. Finally, customer expectations for service experience are rising, informed by customer experience (CX) leaders in other industries such as Amazon, Uber, and digitally native companies, which offer seamless, personalized, and proactive experiences.

Given both the uncertainties and opportunities created by these trends, many leading energy providers are focused on driving CX improvements to boost their resilience. For regulated utilities, this means finding ways to simplify processes, create new digital products and experiences, develop new offerings, and ultimately make it easier for customers to do business with the utility—all in the face of tighter budgets and directives to reduce the cost to serve customers. For utilities in competitive energy-retail markets, this requires optimizing customer experiences and building new capabilities to maximize customer acquisition and retention, while maintaining a competitive cost to serve.

Consider, for example, the case of a North American electric utility. As part of a recent large-scale transformation, it identified several opportunities to modernize customer self-service for subjourneys such as customer bill analysis, making digital payments, and accessing payment assistance. These simpler solutions to customer pain points increased customer satisfaction and reduced call-center operational costs.

However, many energy providers struggle to make real progress on their CX ambitions. A number of CX initiatives stall or ultimately fail to realize their intended impacts, be these financial, customer satisfaction, employee satisfaction, or otherwise.

Typical CX roadblocks in energy

Utilities undergoing CX transformation may encounter challenges. These can include the following:

Ambiguous or imprecise business cases

Unrealistically high expectations for new technology implementations—often accompanied by hockey-stick-like impact curves in the future—make it difficult to rally the organization and secure necessary funding because the link to value seems weak or unclear. Business cases should articulate clearly how the CX changes will improve operational and cost performance and customer satisfaction.

An inability—or unwillingness—to engage deeply with customers

Some organizations find it easier to rely on long-held internal beliefs instead of deeply engaging with the customer (for example, via call listening, interviews, focus groups, or panels). In extreme cases, organizational leaders in the business have never talked with real customers.

The result is experiences designed from the energy-provider's perspective, rather than the customer's. Symptoms of this can include communicating with customers—on the web, phone, or in print—in language that does not resonate with customers but is utility- or regulation-focused (for example, "ratepayers," "service territory," "tariffs").

Or—if organizations do understand the customer's perspective—teams may still overindex on "shiny objects" (such as chatbots) rather than root-cause and design-led solutions. Teams often overrely on core technology changes that take a long time to implement and rarely deliver the expected value. This can be at the expense of investment in the core experience. A broader pitfall occurs when organizations focus on spot fixes or incremental changes rather than on end-to-end transformation.

Doubts about tech ability and agility

In large energy providers, there can be a misconception that legacy technology is a barrier to change. The perception is that antiquated tech stacks and legacy "tech debt" resulting from decades of underinvestment or historical budget cutting will make rapid, modular, and frequent improvements difficult to execute.

A typical manifestation of this, given the silos that can exist between customer and IT teams, is a limited understanding of feasible types of customer experiences and changes. Some organizations believe that only large-scale, long-term replacements hold the answer. The resulting long project timelines—often multiyear, technology-led transformations—rely on "big bang" releases that make it difficult to maintain the required end-to-end focus on CX and operational performance improvements and financial-value delivery.

However, perception is not always reality. Progress is possible and innovative energy providers find ways to deliver CX transformation quickly and effectively. By way of example, one utility was able to make significant changes to its digital CX in nine months, despite having 30-year-old customer-information systems and even older support systems.

Lack of cross-functional collaboration
Functional silos based on ownership of individual customer touchpoints can prevent true, end-to-end customer-journey optimization. For example, multiple US utilities have treated "payment" and "payment assistance" as distinctly separate customer experiences because of existing business processes. However, customers naturally think of these as intricately linked and expect integrated support—such as one customer profile for both.

Despite these challenges, a number of energy providers have "cracked the code" for delivering superior CX, while at the same time reducing overall costs to serve customers. These leaders are delivering real impact in a matter of months, rather than years. And in such time frames, they've increased customer satisfaction by 10 to 20 percent and reduced cost to serve by 20 to 30 percent.

Key lessons can be learned from these organizations on how to achieve such benefits—with an approach that combines new collaboration models, design-led thinking, and modern, agile means of digital delivery to get to results faster.

Three key factors can help to get this right:[1]

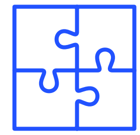

1. *Align on the "why"*
   The first step in a successful digital CX transformation is to align the organization around a common CX aspiration (for example, digitally enabled self-service), informed by the brand promise, to enable ongoing collaboration.

   In particular, leadership alignment on customer importance is critical to drive a shift in culture. Leaders can drive collaboration and alignment between the business and the digital IT teams on the end-state vision. Typically, in energy providers, the CX team focuses on the call center but wants to expand into increasingly digital channels and tech-enabled experiences, while the digital team lacks the customer insight and frontline strategy to optimize digital CX. To develop a common purpose, organizations should bring together the business and technology teams to simultaneously tackle business-process changes and technology changes. The changes can then be delivered through the primary—and increasingly digital—customer channels.

   For example, in a recent CX transformation in a different utility in North America, the organization agreed to target industry-leading CX consistently across all its processes. After creating a collaborative team that merged customer and digital experiences, it is now targeting a reduction of more than 20 percent in its cost to serve customers in some of its core journeys.

   Leadership can then prioritize the journeys and the channels that matter most and establish

[1] Victoria Bough, Ralph Breuer, Nicolas Maechler, and Kelly Ungerman, "The three building blocks of successful customer-experience transformations," McKinsey, October 27, 2020.

a clear link to business and customer value. For energy providers, this likely means that payment and payment-assistance-related experiences—for example, the "I have issues paying my bill" customer journey—are strong starting points because of the large number of customer touchpoints. By way of illustration, at a central US utility, it was observed that 5 percent of customers drove 60 percent of the call volume for these journeys, creating a solid business case for prioritization given the high call-center volumes.

Collections and field activities, such as notices, connects, and disconnects, as well as select back-office activities, can also typically be encapsulated in the scope of these efforts. Conversely, focus on low-utilization channels and applications should be minimized. We repeatedly observe channels that have low utilization relative to their investment levels such as chatbots, smart homes, and native-mobile applications. These are typically not well utilized but carry high development and maintenance costs.

Finally, a clear road map—one that starts with a clear and specific aspiration that is both concrete and achievable—is critical for making progress and tracking consistent. Subsequently, one to two customer journeys should remain in focus at any given time, for end-to-end optimization.

These optimization efforts often include front-end CX changes, back-end process changes, and technology changes, and should be prioritized based on the financial, satisfaction, and strategic value at stake. Each journey can be offered with a consistent approach that diagnoses customer pain points, designs the desired future state, prioritizes changes based on value and feasibility, and plans implementation using an iterative, staged approach that delivers value quickly.

2. *Lead with user-centric innovation to transform the business and optimize customer journeys*
   When mapping out the future-state customer journey, combining traditional CX practices with user-centered design processes has proved to bring great value. Traditionally, design thinking is an approach to solving problems that puts the customers' needs above all else. CX design also uses a human-centered approach[2] but in addition aligns with broader business objectives by establishing the link to value across the entire customer journey. Redesigning customer experiences needs to be grounded in the voice of the customer and the underlying data rather than long-held organizational beliefs.

   For example, rather than an energy provider focusing on chatbots and text-to-pay functionality that do not typically address real pain points, they could instead focus on how best to guide the 5 percent of customers who may drive 50 percent or more of payment-arrangement call volumes to digital channels with new features and customer-friendly language. Consider, for example, a web- and mobile-friendly sign-up tool that utilizes a "quiz" approach to suggest the best payment agreement based on each customer's unique situation.

   Design principles should be included from the beginning to eliminate complexity and to focus the organization on what matters most to customers. For example, rather than breaking down payment arrangements and human support service into different customer journeys, organizations could place all options in one place—which is how customers prefer it.

[2] "CX without design only gets you halfway," McKinsey, June 21, 2022.

Rather than designing websites and customer experiences using utility jargon (“service territories,” “tariffs,” and the like), organizations can use terms that are familiar to customers to ease the customer journey experience.

Supportive CX changes include account and bill simplification, such as a simple, modern cover page on top of pages of required regulatory legalese. Another positive CX change is shifting functionality online—for example, human services. Or the interactive voice response could be simplified, using customer-friendly language to contain calls. These can create a more supportive experience at a time when customers may be experiencing shock and frustration and need to identify their next best actions quickly. The emotional and tactical support not only improves the customer experience but also empowers customer self-service—which can result in improved customer affordability in the long run.

Of course, the call center remains a key driver of CX, and call-center operations must be evolved in parallel to modernizing the CX. While much of the work to modernize CX will occur in digital channels and experiences, delivering enduring performance improvement to both customer satisfaction and cost requires that companies transform how they run their call centers. This typically includes better performance management, training, and coaching of the call-center employees, as well as evolving the workforce strategy and agent-skilling model to handle more complex calls. Operational excellence across the full (and potentially smaller) footprint is also vital.

3. *Execute technology upgrades through new ways of working*
   CX leaders should address technology alongside the first digital product releases. Too often, energy providers undertake massive, multiyear technology upgrades that last years before any value is delivered. Technology changes addressed with the first releases can build momentum and deliver value quickly (ideally in the first six months). For example, engineering practices—such as feature toggles, infrastructure as code, and continuous integration and continuous delivery that accelerate delivery and allow for iterative testing and learning—can help lead to success.

   Such engineering excellence should be supported with a modern CX tech stack and an agile product-based delivery method. For instance, in an East Coast utility’s recent digital CX transformation, a phased upgrade plan was created and implemented alongside each release of new features. This helped to bust the organizational myth of tech stack modernization as a barrier to improving CX.

   However, it is not enough to only build new features and channels and expect customers to alter their behavior on their own. Customer conversion and behavior change is an ongoing, active process that requires careful planning, targeted messaging, and dedicated campaigns to push behavior change and ultimately realize value.

New collaboration models and agile principles help to drive success, which demands alignment among an organization’s customer team, tech team (including digital and IT), and design team. This collaboration should then deliver using agile principles, including iterative design, product teams rather than waterfall projects, biweekly sprints that deliver incremental value at each sprint review, feature backlogs, and full transparency of progress and priorities.

For example, at a leading European energy player, the customer-centric transformation included the setup of cross-functional “excellence rooms,” bringing together stakeholders from commercial functions, customer service, and IT to work in a data-driven, agile setup as the new normal of day-to-day collaboration. The new operating model drastically increased speed of decision making and broke down previous functional silos, leading to a step change in customer-journey excellence and realizing cost-to-serve potentials.

Finally, organizations may need to hire talented individuals who bring skills that may be currently missing in the organization. These include both tech and design skills that do not exist in many energy providers at this stage. Design capabilities go far beyond the historical focus on user experience and often require new talent to build the type of organizational muscle needed in CX. New tech talent may include a host of specialists such as developers, designers, scrum masters, and more. These people will quickly become part of the core capability that the organization will leverage and deploy on an ongoing basis to enhance CX.[3]

These are unprecedented times for energy providers globally. The volatility of market forces in the energy sector are likely to continue in the near future. Digital CX transformation offers an opportunity for the industry to become nimbler and more customer-centric, and energy providers that undertake such efforts will better position themselves to boost their resilience in the face of these market forces.

Each utility and energy retailer will have a different starting point on this digital CX journey. For some, the first step in a digital transformation may be to understand better their starting point and which customer journeys matter most to their customers, and therefore where to focus transformation efforts. For others, it may begin with a technology and agile-maturity assessment to identify both limitations and strengths, or inform prioritization.

For others, it may require hiring new design or digital talent to build the internal capabilities required to deliver on an ongoing basis. Regardless of the starting point, all power providers should maintain a laser focus on the call center to ensure that the call-center CX is optimal. To keep customers satisfied, supply needs to match demand, and call-center operators need to be trained to handle the more difficult calls.

Forward-looking energy providers focus on customer experience not only as a means to boost resilience to challenging industry conditions but also to improve delivery and value capture. Focusing on three core factors when executing digital CX transformation could generate enough cost savings to pay for itself—setting off a reinforcing cycle of evergreen CX optimization to capitalize on the opportunities that lie ahead.

[3] Oliver Ehrlich, Harald Fanderl, David Malfara, and Divya Mittagunta, “How the operating model can unlock the full power of customer experience,” McKinsey, June 28, 2022.

Bobby Dean is an associate partner in McKinsey’s Washington, DC, office, where **Scott Perl** is a partner; **Oliver Ehrlich** is a partner in the Düsseldorf office; and **Hongqiao Lu** is an associate partner in the New York office.

Digital transformations in energy retail: A shift toward advanced platforms

More energy retailers are transforming their digital platforms to cut costs and to stay competitive. We explore approaches to platform transformation and key questions to consider before embarking on a transformation journey.

This article is a collaborative effort by Dago Diedrich, Christoph Fuchs, Paul Küderli, Stephan Mühlhäuser, and Kai Vollhardt, representing views from McKinsey's Electric Power & Natural Gas Practice.

Energy retailers face a pivotal moment. Energy supplies have been at risk in many markets, and prices have gone up. The economic environment has prompted retailers to seek ways of streamlining their operating models and reducing their cost to serve. Meanwhile, customer preferences are moving in favor of renewable energy and new services such as electric-vehicle charging. In response, many energy retailers have been looking for technology solutions that can help improve profitability and customer satisfaction, because outdated operating models and core technologies could be holding some organizations back.

One approach we're seeing more of involves installing a new technology platform—that is, the underlying software to operate the business. This type of digital transformation is also known as replatforming. Retailers implement these new platforms either by partnering with a platform provider or developing the platforms in-house. In some cases, the retailer also implements a new operating model during the same process.

On the front end, the new platforms might offer customers more self-service options in an easy-to-use interface. On the back end, the platforms are often cloud-based, highly scalable, and data-driven. They can also help streamline certain processes, for example, linking customer contracts from various product lines or having core information on customer interactions from different systems readily available for agents.

Replatforming can be a complex process that often involves an overhaul of IT architectures, disruptions to certain day-to-day operations, and a culture shift within the organization. There are different approaches to replatforming, depending on the retailer's current situation (for example, geographic market, product and service offerings, and existing capabilities). Some approaches are more extensive than others and could include a revamp of the whole organization's operating model. And the marketplace for energy retailers is moving at such a fast pace that such a transformation could feel daunting.

The rewards, however, could be significant. Based on our experience with various energy retailers who have already implemented new platforms, organizations have an opportunity to capture significant amounts of market share, improve customer satisfaction, and show profitability.

In this article we'll explore the upside of replatforming, markets where replatforming is having an impact, three approaches to replatforming that organizations can take, and key questions energy retailers should consider before embarking on a platform transformation.

The upside of replatforming

Based on our experience, companies that undergo successful replatforming transformations often see three key benefits.

Increased customer satisfaction. Organizations are able to offer higher service levels with significantly higher average speed of answer (ASA) to customer inquiries, improved digital self-service offerings, and more personalized interactions through easier access to customer data. We have seen customer satisfaction improvements of up to 300 percent in some cases.[1]

Decreased cost to serve. The new platforms can make customer operations more efficient for employees by, for example, making all required information directly accessible to the agent, including background from the customer's previous interactions and a complete overview of the customer's contract status. In some cases, we have seen a 50 percent decrease in cost to serve.

Increased speed to market for new products. Organizations have the flexibility to react to market developments and changing customer preferences by quickly incorporating new products and service offerings into the platform. We have seen a 50 percent increase in time to market for new products.

[1] Measured on third-party platforms such as Trustpilot.

Replatforming is already having an impact in many markets

Energy retailers operate in markets where regulators allow businesses to compete for gas and electricity customers—for example, in Australia, Germany, and the United Kingdom. (In other markets, customers might purchase their energy from a single utility.) In these competitive markets, particularly ones where energy prices have increased, platform transformations are often under way or in the planning stages. Some of these markets provide a glimpse of how new players have taken market share, and how being early to market with a new technology can potentially provide a competitive advantage.

Australia, Germany, and the United Kingdom are regions where many energy retailers are relatively ahead of the curve on replatforming, closely followed by markets such as France and Japan, where several incumbents are already planning larger replatforming efforts. A few regional examples include the following:

— *United Kingdom.* Retailers including E.ON, British Gas, and SSE have either partnered with platform providers or developed new platforms.[2]

— *Germany.* Retailers including E.ON and EnBW have partnered with platform providers.[3]

— *Australia.* Major retailers have either switched to a new platform or are considering doing so.

Market disruption

One key observation we've made in certain markets is that a new player builds an advanced platform and then potentially enters the market as both a retailer and a software provider. The new player grows on multiple fronts: building a customer base with its retail business and selling its software to incumbents, which can run the technology off the new player's platform. We've seen the following scenario play out in a few cases:

— *New market entry.* A new player enters the market with an advanced proprietary platform; the new platform builds an initial customer base.

— *Scaling and first partnerships.* The new player scales the platform and grows, often by acquiring smaller incumbent players or forming joint ventures. The new player offers initial partners exclusivity on the platform for a certain time frame.

— *Other incumbents invited.* After the initial period of exclusivity, other incumbents are invited to move to the platform.

— *Necessity for survival.* Remaining market players that haven't moved to the new platform recognize that modernizing their platform could be needed to remain competitive.

This scenario shows that incumbents that are late to replatforming could not only lose market share but also have a more limited course of action. For example, late-coming incumbents might have to wait out the exclusivity period of the new player's platform or explore an alternative platform provider that is potentially less advanced in its technology.

Based on these experiences, we believe it is key for incumbents to be ahead of the replatforming curve. Quick movers to replatforming could be setting themselves up for the following advantages:

— *Accelerating future growth* by reducing cost to serve—and potentially keeping prices more competitive for customers.

— *Decreasing customer churn* by improving customer service and experience relative to competitors' outdated platforms.

— *Improving strategic options* by having the opportunity to partner with software providers during exclusivity periods.

[2] "E.ONnext – E.ON and Kraken Technologies form strategic partnership for E.ON's UK residential and commercial customer business," E.ON press release, March 23, 2020; Emma Lovell, "Centrica: Building a new British Gas with ENSEK," ENSEK, April 26, 2022; "OVO and SSE: Questions about OVO Energy, SSE and Kaluza?" Kaluza, accessed November 16, 2022.
[3] "References," powercloud, accessed November 16, 2022.

The sooner a retailer starts to recognize developments in the market and consider a course for its platform, the better chance it will have of remaining competitive. In our experience, players that are early movers to replatforming end up having more time to strengthen both their platforms and their market positions than players that follow later.

Three common approaches to replatforming

We have seen, in our experience, three primary archetypes for replatforming transitions. Each approach comes with a set of potential advantages and considerations based on the retailer's current platform and capabilities, geographic market, and product offerings.

— *Selective transformation of some processes or technical capabilities.* This approach improves certain aspects of the operating model and technology backbone but leaves in place functions that are already successful. For some companies, selective transformation is a predecessor to a more extensive replatforming process. For example, a German energy player opted for a selective transformation as a way to better integrate existing core elements (such as the product definition engine) with new technology capabilities for customer service operations. In this transformation case, the company saw considerable improvement in earnings before interest and taxes (EBIT) with relatively limited organizational effort.

— *Platform replacement.* This approach entails complete retirement of the current platform and migration to a new platform—in other words, the back-end technology is completely overhauled and modernized. In the German market, for example, some players have chosen platform replacement in scenarios where the core technology of a platform was the key factor holding back potential efficiency gains.

— *A holistic transformation of technology, operating model, and organization.* This approach requires not only platform replacement but also significant changes to an organization's entire operating model and, in some cases, its business (offering new products or more self-service options to customers, for example). Processes apart from customer service and operations, including HR and finance, could be transformed along with the technology. In such a case, an unprofitable incumbent that's losing a significant amount of money every week might decide that a holistic transformation provides the best chance of returning to profitability. Multiple UK energy retailers have successfully conducted holistic transformations, which typically entail a platform redesign from scratch; significant changes to the operating model; a lean approach to support functions; a zero-based[4] approach to the technology stack (getting rid of unnecessary legacy software, for example); and the introduction of a new company culture—one that is nimble and able to embrace change.

Recent cases tend to show that holistic replatforming transformations are more successful than platform replacements in delivering sustainable impact and driving overall value. (For more on holistic transformations in the United Kingdom, see sidebar, "How one UK player executed a successful holistic transformation.") Platform replacement alone can require great effort to adapt existing business processes and may not maximize the full potential of a new technology. But platform replacement could suit an organization with limited appetite for change but a pressing need to overhaul outdated technology. A selective transformation is most narrow in scope but could be a viable option for an organization that's limited in its funds or capabilities. While it's true that the more extensive approaches (platform replacement or holistic transformation) can be more difficult to execute, we find that over the long term, a new platform—if done right—has a higher potential to deliver cost savings and an improved operating model.

All three archetypes likely require a significant change to existing IT architectures, day-to-day operations, and workplace culture. A digital

[4] For more on zero-based strategies, see Søren Fritzen, Matt Jochim, Carey Mignerey, and Mita Sen, "Zero-based productivity: The power of informed choices," McKinsey, July 25, 2018.

transformation could take 12 to 24 months and might require reskilling talent, bringing on new tech talent, and working with agility. A transformation that's not executed successfully could risk a loss of customers and employees, as well as damage to the brand.

Based on our experience, successful organizations ensure that a set of prerequisites is in place at the beginning of a transformation. These organizations typically establish a simple transformation governance structure to streamline decision making (appointing a single transformation leader, for example), adopt agile ways of working and anticipate fast iteration cycles, and set a standard for clear and transparent communication throughout the transformation.

Key questions for energy retailers to consider

Energy retailers that are looking to replatform could start the journey by exploring a set of key considerations. Retailers can work to understand the current state of their local market, with rigorous observation of the market players' activities; decide which transformation approach would be most suitable; and build a solid implementation plan. The following questions can help to guide the decision for replatforming activities and ensure that the decision is not only technically driven but also business driven:

— *For understanding the current state of the local market:* To what degree have competitors begun replatforming? Have competitors made relevant technology acquisitions or partnerships? What are the risks of *not* undertaking a transformation in the current marketplace (for example, potential loss of customers to competitors)? Understanding the local market will help inform how quickly a company might want to move on a transformation and what partnerships with software providers are possible.

— *For considering which replatforming approach is most suitable:* Which archetype of transformation fits the company's future goals? What is the current capability to execute

How one UK player executed a successful holistic transformation

E.ON Next, part of the UK division of the E.ON Group, is an energy retailer that serves more than six million customers. Following a merger with NPower as part of a larger corporate deal, E.ON UK experienced heavy losses. In 2019, the organization partnered with a platform provider to execute a holistic platform transformation. E.ON UK revamped its technology platform for customers, as well as its operating model, and rebranded the UK retail energy business E.ON Next. The transformation, by the company's measures, has been an overwhelming success, leading to improved customer satisfaction and increased profits. Leadership at E.ON Next shared these lessons about the experience:

— *View the platform provider as a partner.* The platform provider acted as more than a supplier of technology. It actively helped in building a new operating model and fostering a more agile company culture. Leadership at E.ON was also open to the provider continually challenging both the transformation approach and previous processes as a way to keep ambitions high.

— *Give key decision makers autonomy to deliver.* The transformation had a ring-fenced leadership team that had autonomy within a clear framework and required involvement from the wider business only when there were deviations from that framework.

— *Make transformation priorities simple and clear.* Setting a clear direction helped keep teams working toward a common goal. An example of one guideline: "The only thing that counts is to create the best possible customer experience."

a transformation? This decision could ultimately come down to the current financial pressure on the organization.

— *For determining whether to "build or buy" a new platform:* Which platform providers are available in the market? Which option is likely to deliver greater impact across the organization? One factor leaders can consider is the amount of tech debt the company has already incurred—that is, the work a company needs to do to update its IT landscape. A lower tech debt could indicate that an organization has a better starting position for building in-house, whereas a high tech debt might make a platform provider a more attractive option.

— *For creating an implementation plan:* How can the entire organization come on board and help carry out the transformation? What are critical factors that can guide the transformation? What are lessons learned from previous transformation projects that can guide the next transformation? How does the organization ensure the transformation achieves value expectations? A starting point for an implementation plan could include reaching out to an energy retailer in another geography that has successfully replatformed to better understand best practices and potential pitfalls.

Energy retailers in many markets face margin pressure, new customer demands, and increasingly stiff competition. Embarking on a platform transformation is one way retailers can stay ahead in the marketplace. Replatforming is a significant undertaking, but proactive players could be looking at increased customer satisfaction, a more efficient operating model, and more flexibility to change with the times.

Dago Diedrich is a senior partner in McKinsey's Düsseldorf office; **Christoph Fuchs** is an associate partner in the Munich office; **Paul Küderli** is a partner in the Frankfurt office, where **Stephan Mühlhäuser** is an associate partner and **Kai Vollhardt** is a senior partner.

The authors wish to thank Christoph Heilmann and Mona Zimmermann for their contributions to this article.

An AI power play: Fueling the next wave of innovation in the energy sector

This article explores how Vistra Corp. is partnering with McKinsey to improve efficiency and reduce emissions by using AI.

by Barry Boswell, Sean Buckley, Ben Elliott, Matias Melero, and Micah Smith

Matthias Radochla / EyeEm

Tatum, Texas, might not seem like the most obvious place for a revolution in AI, but in October 2020, that's exactly what happened. That was when Wayne Brown, the operations manager at the Vistra-owned Martin Lake Power Plant, built and deployed a heat rate optimizer (HRO).

The "heat rate" is a measure of the thermal efficiency of the plant—the amount of fuel required for each unit of electricity produced. To reach the optimal heat rate, plant operators continuously monitor and tune hundreds of variables, or "set points" on things such as steam temperatures, pressures, oxygen levels, and fan speeds.

It's a lot for any operator to get right 100 percent of the time—so Vistra thought AI could help.

Partnering for innovation

With this goal in mind, Wayne and his group worked together with a McKinsey team that included data scientists and machine learning engineers from QuantumBlack, AI by McKinsey, to build a multilayered neural-network model—essentially an algorithm powered by AI that learns about the effects of complex nonlinear relationships.

This model went through two years' worth of data at the plant and learned which combination of external factors—such as temperature and humidity—and internal decisions, such as set points that operators control, would optimize the algorithm and attain the best heat-rate efficiency at any point in time.

Vistra team members provided continuous guidance about the intricacies of how the plant worked and identified critical data sources from sensors, which helped McKinsey engineers refine the model, adding and removing variables to see how those changes affected the heat rate.

Through this training process and by introducing better data, the models "learned" to make increasingly accurate predictions. When the models were accurate to 99 percent or higher and run through a rigorous set of real-world tests, a McKinsey team of machine learning engineers converted them into an AI-powered engine. This generated recommendations every 30 minutes for operators to improve the plant's heat-rate efficiency.

At a meeting with all of Vistra's leaders to review the HRO, Lloyd Hughes, an experienced operations manager at the company's Odessa plant, said,

> **'There are things that took me 20 years to learn about these power plants. This model learned them in an afternoon.'**

–Lloyd Hughes, operations manager

Sidebar

About Vistra and its emissions reduction goals

Vistra Corp. is the largest competitive power producer in the United States and operates power plants in 12 states with a total capacity of more than 39,000 megawatts of electricity—enough to power nearly 20 million homes.

Vistra has committed to reducing emissions by 60 percent by 2030 (against a 2010 baseline) and achieving net-zero emissions by 2050. To achieve its goals, the business is increasing efficiency in all its power plants and transforming its generation fleet by retiring coal plants and investing in solar and battery energy storage, which includes the world's largest grid-scale facility for battery energy storage. Vistra's to path to net zero will also require it to grow its zero-carbon portfolio to more than 7,000 megawatts by 2026.

"There are things that took me 20 years to learn about these power plants. This model learned them in an afternoon."

With this kind of power at their fingertips, Wayne and his team could make better, more informed decisions. Acting on the HRO recommendations helped Martin Lake run more than 2 percent more efficiently after just three months in operation, resulting in $4.5 million per year in savings and 340,000 tons of carbon abated. This carbon reduction was the equivalent of taking 66,000 cars off the road.[1] If that doesn't sound like a lot, consider this: companies that build gas-fueled power plants invest millions of dollars in research and development over four to five years to achieve 1 percent improvement in power generation efficiency. Vistra hit that improvement level in only *one-twentieth* the amount of time, using the data and equipment it already had.[2]

Vistra has since rolled the HRO out to another 67 power generation units across 26 plants, for an average 1 percent improvement in efficiency and more than $23 million in savings. Along with the other AI initiatives, these efforts have helped Vistra abate about 1.6 million tons of carbon per year, which is 10 percent of its remaining 2030 carbon reduction commitment. That's equivalent to offsetting about 50 percent of what a 500-megawatt coal plant emits.

What happened at Martin Lake has happened at dozens of Vistra's other power plants, with more than 400 AI models (and counting) deployed across the company's fleet to help operators make even better decisions. It also reflects a core trait of Vistra's AI transformation, which is that it isn't a story of one massive hit, but rather the story of dozens of meaningful improvements snowballing to deliver significant value in terms of accelerating sustainable and inclusive growth. It's also the story of how an organization created an approach to rapidly scale every successful AI solution across the entire business. And it's a story of how a culture of continuous improvement, combined with a powerful AI modeling capability, helped leaders and plant operators do their jobs better than ever.

[1] Calculations based on "Greenhouse Gases Equivalencies Calculator - calculations and references," US Environmental Protection Agency, last updated May 30, 2023.
[2] Calculations based on source material from "A Brief History of GE Gas Turbines," Power, July 8, 2019, and Frequently asked questions, US Energy Information Administration, last reviewed September 20, 2021.

Sidebar

What does 'machine learning operations' mean?

Machine learning operations is the set of practices and infrastructure that manages the production and deployment of AI solutions or products. Improvements in AI tooling and technologies have dramatically transformed AI workflows, expedited the AI application life cycle, and enabled consistent and reliable scaling of AI across business domains. This framework enables organizations to create a standard, company-wide AI "factory" capable of achieving scale.

With more than $60 million captured in about one year of work and another $40 million in progress, Vistra is well on its way to delivering against a road map of $250 million to $300 million in identified EBITDA and more than two million tons of carbon abatement per year. The AI-driven advances at Vistra have heralded a generational shift in the power sector in terms of improvements in efficiency, reliability, safety, and sustainability.

If the 1 percent efficiency improvement the HRO delivered across the fleet were carried across all coal- and gas-fired plants in the US electric-power generation industry, 15 million tons of carbon would be abated annually—the equivalent of decommissioning more than two large coal plants or planting about 37 million trees.[3] That means less fuel needed to deliver power to the hospitals, schools, and businesses that rely on it. AI has the potential to bring similar levels of improvement to renewables as well, making them a more cost-effective and attractive energy option.

Turning points on the AI journey

Healthy skepticism and a culture of favoring action over words at Vistra meant that the biggest hurdle in the AI journey wasn't the technology—it was the people. Vistra leadership and operations managers needed to know what AI could do and be convinced it could really work.

It was this ingrained culture of continuous improvement—alongside a highly competitive market and a commitment to sustainability—that convinced Vistra's leadership they needed to give AI a chance.

Seeing real possibilities

The first question was a relatively simple one: "How can AI help improve the way Vistra generates power?"

Answers to that question bubbled over when 50 of Vistra's top leaders came to a McKinsey-hosted workshop. In multiple sessions, experts explained how AI worked, walked through detailed case studies showing how other companies used analytics and AI to generate value, and gave live demonstrations of technologies, including digitalized workflows and machine learning. Leaders in analytics from various sectors—including Amazon, Falkonry, Element Analytics, and QuantumBlack, AI by McKinsey, as well as G2VP from the venture capital world—provided insights and examples of how AI works.

[3] The tree equivalency is based on calculations derived from Jean-Francois Bastin et al., "The global tree restoration potential," *Science*, July 5, 2019, Volume 365, Number 6448.

Exhibit

Four types of AI solutions contributed to an annual $61,355,000 impact.

Capacity forecasting

Predict the true capability of the plant for total electricity generation based on ambient conditions, plant conditions, and wear and tear

9 AI model deployments

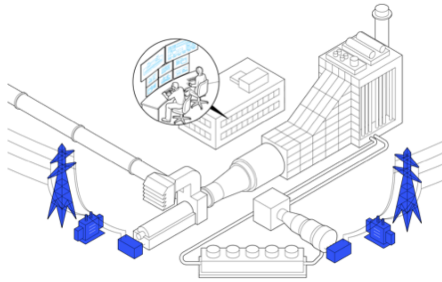

Transmission and distribution to customers

Implementations were performed on both combined-cycle gas turbine (CCGT) (shown) and coal-fueled plants and are currently progressing to renewable solar and battery storage systems.

Plant reliability

Predictive maintenance—Use advanced analytics to monitor equipment performance and predict disruptions or failures with enough lead time to take corrective action

120 AI model deployments

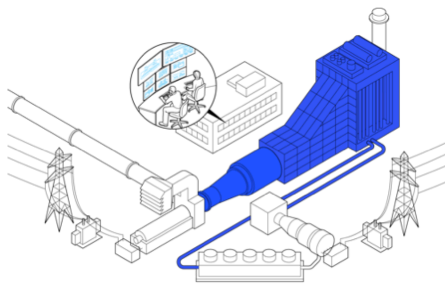

Heat recovery steam generator health, gas turbine health and temperature spread, pump health

Values-based maintenance—Use advanced analytics to find the optimal trade-off between the time and cost of repair work and the corresponding improvement in plant performance

130 AI model deployments

Overall annual carbon abatement 430,000 tons

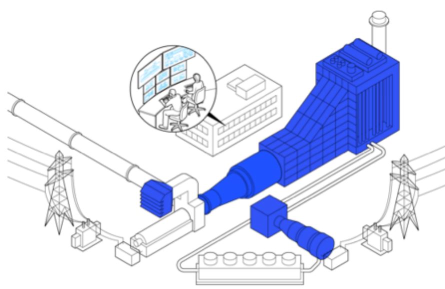

Heat recovery steam generator cleaning, compressor washes, air inlet filter replacements, condenser cleaning, generator windage losses

Implementations were performed on CCGT (shown) and pulverizers and scrubbers in coal-fueled plants and are currently progressing to inverters for solar and battery storage systems.

Efficiency (heat rate solution)

Optimize the amount of electricity generated for each unit of fuel consumed

60 AI model deployments

Overall annual carbon abatement 1,207,000 tons

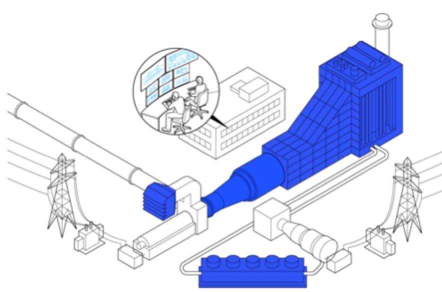

Steam temperatures and duct firing, operating mode, inlet cooling, cooling fan operation

Implementations were performed on both CCGT (shown) and coal-fueled plants and are currently progressing to renewable solar and battery storage systems.

Performance management

A series of tools to help optimize various aspects of plant operations on a daily basis

399+ AI model deployments

Overall annual carbon abatement 1,637,000 tons

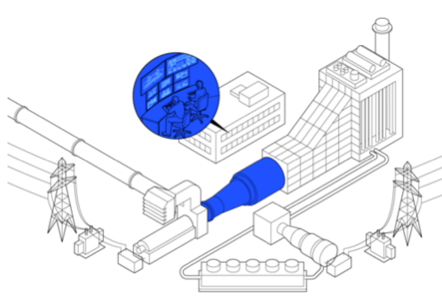

Optimize start-up time

Implementations were performed on CCGT (shown). Optimizing auxiliary load,optimizing pump and fan operations,-detecting and troubleshooting equipment anomalies, and improving controls of furnace balance, coal feed, and limestone additions were performed in coal-fueled plants. Performance management tools are under development for solar and battery installations.

"I saw an example of how a metallurgical plant was using AI to help its operators optimize set points and it clicked for me," says Patrick "Cade" Hay, the plant manager for Vistra's Lamar power plant. "I saw how I could translate that into helping me run my own plants more efficiently. This was my light bulb moment."

Company leaders and plant managers pored over process flow sheets and engineering diagrams to determine pain points as well as opportunities. This exercise allowed them to focus first on finding where the value was, then on what technologies were needed to deliver it. Many of the operations opportunities were around yield and energy optimization and predictive maintenance, which, according to our research, were the top AI use cases for manufacturing industries.[4]

By the end of the session, Vistra had developed a strategy to develop a series of AI solutions that could capture $250 million to $300 million in potential EBITDA while helping the company achieve its 2030 emissions reductions goals.

Codeveloping the AI

While the analysis looked promising, proving it in the field was what mattered. "If our plant managers aren't bought in, then things don't happen," explained Barry Boswell, Vistra's executive vice president of power generation operations. "So we said, 'Let's pick a leader who is knowledgeable and skeptical, because if we can win them over, we can get everyone.'"

They picked Cade. Not only did he run a top-performing plant in terms of profitability, reliability, and heat rate, but he also had a reputation for telling it like it is—Barry trusted Cade to tell him whether or not the value potential in AI was real. When he approached Cade about testing a proof of concept to optimize duct burners, he was intrigued but predictably skeptical. Cade saw the potential in AI but was interested in finding out if it could actually help in the field.

Duct burners essentially work like afterburners in jet planes; they provide a surge of energy when needed. Operators use them as supplements to hit energy targets, which are known as their "dispatch point." The issue is that powering duct burners uses more fuel than regular methods, so it's more expensive, generates more carbon emissions, and increases the wear and tear on equipment.

McKinsey subject matter experts, data scientists, and analytics translators from QuantumBlack, AI by McKinsey, worked closely with a Vistra team made up of power generation and process experts as well as frontline operators to understand how the plant works, what data was available from the sensors already in place, and what variables could be directed—such as the fact that the number of cooling fans running could be controlled, but the ambient temperature couldn't.

As the teams developed the models, plant operators reviewed recommendations to see what made sense, what other variables needed to be tested, and what kinds of recommendations the operators would find most helpful. By analyzing the effect of various inputs and set points on the plant—such as pressure and humidity, the angle of blades in the gas turbine, usage of inlet cooling technologies, and the age and performance of various components such as filters and condensers—and running it through the model, the analysis was clear: overall duct burner usage could be reduced by approximately 30 percent, which would result in the equivalent of $175,000 in yearly savings on fuel costs and wear and tear on the system, in addition to abating about 4,700 tons of carbon per year.

"We worked closely with the team from McKinsey to develop AI models that reflected the realities of how power plants operate," said Cade, "and then when we saw the recommendations coming out of the AI tool, I saw how much real value there was."

Going for scale and adoption from the beginning

Vistra's leadership realized from the beginning that the only way to achieve their efficiency and carbon abatement aspirations was to scale every solution. "We manage the Vistra fleet as one. If a plant is doing something that works, we want every plant to do it," says Barry. "That's what we're built to do."

[4] "The state of AI in 2021," McKinsey, December 8, 2021.

That realization led Vistra to invest in a five-part system to scale and sustain AI solutions:

1. Turn each successful proof of concept (or MVP) into a product

A standardized solution would need to be deployed at each power plant and would become easier to maintain over time. When a solution has proven value at a pilot site and is approved for scaling, a team of software and machine learning engineers immediately takes over to refactor, modularize, and containerize the code. That way, there is a single software "core" package for each deployment that can be updated and improved. A product owner manages the overall process and takes ownership for use and adoption.

Over time, the team developed seven solution archetypes, which provided consistent approaches, logic, assumptions, and algorithmic elements as a basis for each new application being developed. This gave each new solution a big head start when development began. It took ten to 12 weeks to build the first HRO. Rolling each HRO out to subsequent new plants now takes just two to three weeks.

2. Create machine learning operations infrastructure

Vistra implemented a machine learning approach to essentially create a "factory" that standardized the deployment and maintenance of more than 400 AI models. At a high level, this approach enabled the team to bring live data from each of Vistra's power units into a single database; use GitLab software to manage version control for code; containerize the code so it could be easily deployed to any environment; set up a scheduler—using Apache Airflow—to make sure recommendations and actions were delivered on time; create dashboards to monitor model performance and usage; and

Sidebar

Working together

Denese Ray, operations shift supervisor for Vistra's Coleto Creek power plant; Doug Richter, the shift supervisor at the same plant; and Muro Kaku, McKinsey analyst and data scientist, discuss how they worked together.

Denese: "Every week, I'd bring in a different supervisor to join us for our check-in calls with McKinsey. In this way, Muro and the McKinsey team saw what we were dealing with. We really needed all our supervisors and their feedback to build the model so we could get the most from it."

Muro: "That really mattered. For example, when we first trialed the HRO at the Coleto Creek plant, the recommendations were frequently rejected by the operators. We didn't know why. When we spoke with the operators, they told us about specific rejection reasons: for example, the tool recommended increasing superheat temperature by 30 degrees in a 15-minute interval to maximize the heat rate, but it takes the plant longer to heat up, so they rejected the recommendation. We then added a constraint to the tool to keep the increase and decrease to five degrees in a 15-minute interval, which operators could achieve."

Doug: "During load changes [increases or decreases in power output], the parameters are changing constantly, so recommendations to increase air flow at that time were counterproductive. We worked with McKinsey to have the recommendations stop for 30 minutes while we were changing load until the plant was stable again."

Denese: "We were all pretty skeptical of the tool at first, but when we got to play with the heat rate optimizer and see how it worked, and how well McKinsey worked with us in the plant, we understood how it could help."

manage the continuous improvement of each model to make sure the plants were sustaining captured value.

Teams also incorporated multiple approaches to reducing risk by building functional limits, such as maximum throttle pressure or heat levels into the code, and put all code through biweekly peer reviews and multiweek testing. McKinsey's risk dynamics experts worked with the team to test assumptions, review code, and ensure that all risk best practices were reflected in the models.

3. Customize and adapt for 'last mile' implementation

Because each product is designed to be modular and reusable, 50 to 70 percent of each is ready from day one to be used when it's rolled out to a power plant. But customization is always needed because each plant has its own unique characteristics.

Take the maintenance solution developed to understand the best time to replace inlet filters at gas turbines. Each filter costs $150,000, and a plant needs to be shut down for two days to replace one. At the Moss Landing Power Plant in California, the unit that faces the ocean had to deal with a lot more moisture and salt than the one facing inland, so the degradation profile for each filter was different. Similarly, plants in Texas have to manage for dust, while those in Connecticut battle cold weather—local conditions create their own unique degradation profiles.

So dedicated customization teams made up of data scientists, engineers, operators, and power generation experts worked with each power plant to tailor the solution to the unique conditions of that particular plant.

4. Build capabilities

Building on Vistra's well-established culture of continuous improvement, McKinsey worked to help Vistra get the most from AI. That included running a two- to three-week program for operators at each plant to explain the models and how they were developed, as well as instructions on how to use the tool itself.

As Rachit Gupta, Vistra's vice president of generation and wholesale technology applications, said, "People had to know what the model was doing and learn to trust that it was right. Once they saw that the models were generating recommendations that made sense and lowered heat rate, they were ready to start using them more."

Additional training for the core tech team covered how to build and maintain models and resolve issues, as well as deeper machine learning and analytics skills to understand how models work in detail. This happened largely by working side by side with McKinsey AI experts.

5. Design for the operator experience

From the beginning of the solution development process, McKinsey designers worked with operators to understand what their day-to-day activities looked like. What soon became clear was that plant operators had real constraints on their time and had to manage dozens of inputs tracked on an array of screens in the control room. Adding to that workload would be a surefire way to overwhelm operators and reduce the effectiveness of the solutions. The tools had to make operators' lives easier.

For this reason, the screen that displayed the AI solutions and recommendations were integrated into PowerSuite, an interface that operators already used, so they didn't need to monitor yet another screen. The displays themselves were designed to be easy to read. A solution displays a green signal if the plant is running optimally for the given conditions and red when it isn't. A brief recommended action accompanies any red display, with the value attached for implementing that recommendation.

"There's a simple screen that not only shows you what needs to be optimized, but how much it's worth," says Lloyd. "That really registers with people."

Sidebar

AI solution deep dive

QuantumBlack's Ayush Talwar, an expert associate partner, provides more details on the AI model development:

What kind of models did you use?
We used a range of models to fit the specific needs of the plant and the solution. These ranged from Bayesian regression models to deep-learning models.

How did you choose which model to use?
We wanted to find the right balance between model performance, meaning its accuracy; explainability, which describes relationships to the operators; the source richness, which refers to the amount of data available; potential actionability, or how many of the most important features in the model people can actually control; and maintainability, which will reduce the need for more technical skills to maintain them.

How would you compare the models you used with older ones?
Our models were more accurate than previous models and, importantly, more actionable. We used a variety of metrics to measure model performance, such as the mean absolute percentage error (MAPE) and root mean squared error (RMSE). Our models typically achieved MAPEs of less than 1.5 percent.

How did you build up your models?
Each neural-network model was made up of several layers, with each layer containing batch normalization, dropout, and activation. We created cross validation and out-of-sample testing sets using time-based splits to prevent overfitting models. We added features based primarily on domain expertise.

How did you turn models into tools operators could use?
We used machine learning models to make accurate predictions of outcomes. We then wrapped meta-heuristic optimization algorithms around those models to generate recommendations that helped operators make decisions about how to run the plant better. The models were then embedded within the existing production workflows and deployed live in the operator room.

Vistra's story is far from finished. As Cade put it, "We're just looking at the tip of the iceberg." Vistra's road map for 2022 and 2023 includes bringing AI to its rapidly growing renewables fleet of solar and batteries to optimize yield and reliability, among other initiatives.

To help sustain this ambition, Vistra is building up its talent bench. In addition to hiring a small team of data scientists and engineers, Rachit has partnered with the University of Texas at Dallas to offer basic, intermediate, and advanced courses in AI and analytics for Vistra employees. About 70 people have already completed courses, including those reskilling from statistics to machine learning. Vistra has also built relationships with local colleges and universities to develop internship programs and work with students in capstone projects to identify top technical talent.

"We can't sit around and just do what we did yesterday to be ready for tomorrow," Barry says. "We've seen enough to know what's possible."

The authors wish to thank Wayne Brown, Rachit Gupta, Patrick "Cade" Hay, Lloyd Hughes, Denese Ray, and Doug Richter from Vistra Corp., and Richard Bates, Dan Hurley, Pablo Illuzzi, Nephi Johnson, Muro Kaku, Jay Kim, George Lederman, Abhay Prasanna, Noel Ramirez, and Ayush Talwar from McKinsey for their contributions to this article.

The AI-enabled utility: Rewiring to win in the energy transition
September 2023

Cover image:
© Photography by ZhangXun/
Getty Images

Designed by Leff

www.McKinsey.com

@McKinsey
@McKinsey