EET&D
MAGAZINE
Quarterly Issue 3, 2026 — Volume 32

HOW UTILITIES TRACK LARGE CAPITAL PROGRAMS

Photo credit: zzoplanet

EET&D MAGAZINE

Quarterly Issue 3, 2026 — Volume 32

Publisher: Steven Desrochers | steven@electricenergyonline.com
Editor in ChieF: Elisabeth Monaghan | elisabeth@electricenergyonline.com
Graphic Design: Jaguar Media Inc. | tarah@jaguar-media.com

ELECTRIC ENERGY MAGAZINE IS PUBLISHED 4 TIMES A YEAR BY:

JAGUAR EXPO INC: PO Box 50514, Carrefour-Pelletier, Brossard, QC Canada J4X 2V7
Tel.: 888 332.3749 | info@electricenergyonline.com | electricenergyonline.com

POWER POINTS

How AI is Lighting up the Electric Energy Sector

Elisabeth Monaghan, Editor-in-Chief

Today, AI is no longer abstract. In fact, we're seeing everyone from large corporations to small businesses deploy AI to troubleshoot problems, increase productivity, analyze data, write marketing or creative briefs and serve as customer service chatbots. In the electric energy industry, AI has moved from strategy and possibility into a viable presence embedded in field forms, invoice validation, fault detection systems and customer engagement platforms.

THE GRID TRANSFORMATION FORUM

How Utilities Track Large Capital Programs: Field Data Collection to Real-Time Balanced Scorecard at Scale

Hari Vasudevan, KYRO AI, Mythili Chaganti CenterPoint Energy, Jerry Caldwell, CenterPoint Energy, Chris Harris, CenterPoint Energy

CenterPoint Energy's largest grid resilience infrastructure plan to build the most resilient coastal grid in the country is currently underway, with the Public Utilities Commission of Texas approving the more than $3 billion investment. CenterPoint Energy will continue its investments in the electric transmission system, including elevating more than 99% of substations above the 500-year floodplain and upgrading or rebuilding more than 2,200 transmission structures to improve overall reliability and resiliency.

GREEN OVATIONS

The Human Grid Asset: Scaling Grid Stability through Gamification

Alex Hill, Sendero Consulting

While billions of dollars have been invested in smart hardware...grid stability remains at risk because the humans at the end of the wire remain largely unengaged. Getting customers engaged in solutions requires a new approach: activating the behavioral layer of the virtual power plant and cultivating the "human grid asset" - a strategy that turns aggregated demand flexibility into a resource that transforms passive customers into active grid participants.

The Pole Problem: Closing the Field Gap in Wildfire Mitigation Before the Next Season

Davoud Zamani, Gridwap

Undergrounding works. It is also, by the numbers, a generational program. PG&E launched the country's largest such effort in 2021, targeting 10,000 miles. It reached 1,000 completed miles in October 2025, and projects 1,600 by the end of 2026, an estimated 18% reduction in system wildfire risk. That is four years of the most aggressive undergrounding push any U.S. utility has attempted. It is also 16% of the goal.

Building Supply-Chain Resilience in the Age of the Transformer Shortage

James Whaley, GRID Infrastructure Solutions

For years, utilities have invested heavily in grid resiliency: hardening lines, modernizing substations, deploying microgrids and preparing for extreme weather. But now, supply chain resiliency has emerged as an equally essential discipline. If grid resiliency is about keeping the lights on during a storm, supply chain resiliency is about ensuring the equipment needed to build, expand and modernize the grid actually arrives.

Managing AI Within Energy Fleets: Why the Hardest Problems Start After Deployment

Andrew Foster, IOTech

Instead of asking if AI belongs in grid operations, people now talk about where and how to use it. Battery energy storage systems use predictive simulations at the asset level. Solar fleet operators use machine learning to help decide which maintenance tasks to prioritize. Distribution utilities are trying out edge-based analytics for things like voltage optimization, fault detection and demand response across their growing portfolios of distributed energy resources.

Why Half of Utility AI Investments Aren't Paying Off, And What Changes That

Everardo Camacho and Scott Foster, Delta Energy Communications

Across industries, 79% of enterprises have adopted AI agents in some form, but only 11% have them running in production — the largest deployment backlog in enterprise tech history — while a separate cross-study analysis finds 95% of generative AI pilots fail to deliver expected returns. But for utilities specifically, the obstacle is rarely the model. It's that AI tools optimize a single narrow task without ever connecting to the broader operational picture.

Materials as an Asset Management Variable in T&D Reliability

Dr. Nicolas Imlinger, Wacker Chemical Corp.

NERC's 2025 Long-Term Reliability Assessment identifies data centers and other large commercial and industrial loads as the leading source of projected demand growth over the next decade. Equipment manufacturers and utilities often evaluate these demands through electrical ratings, mechanical loads and service conditions. Material chemistry sits underneath each requirement because the formulation and processing of an insulating elastomer determine how the finished component behaves over time.

EET&D MAGAZINE

Quarterly Issue 3, 2026 — Volume 32

Publisher: Steven Desrochers | steven@electricenergyonline.com
Editor in ChieF: Elisabeth Monaghan | elisabeth@electricenergyonline.com
Graphic Design: Jaguar Media Inc. | tarah@jaguar-media.com

ELECTRIC ENERGY MAGAZINE IS PUBLISHED 4 TIMES A YEAR BY:

JAGUAR EXPO INC: PO Box 50514, Carrefour-Pelletier, Brossard, QC Canada J4X 2V7
Tel.: 888 332.3749 | info@electricenergyonline.com | **electricenergyonline.com**

GUEST EDITORIAL

Separating Signal from Noise: How Vertical AI Drives High-Definition Customer Engagement

Gautam Aggarwal, Bidgely

Today, vertical AI is helping energy providers modernize legacy platforms to keep modern customers engaged and satisfied. Unlike generic models trained on broad internet text, energy-specific AI is engineered to separate signal from noise in massive streams of smart meter data. It filters out grid fluctuations to isolate the precise energy “signatures” of individual appliances, such as HVAC cycles, electric vehicle (EV) chargers, or water heaters.

GUEST EDITORIAL

AI Energy Trade-Off: Refining Grids for a Greener Tomorrow

April Miller, ReHack

Artificial intelligence is creating new opportunities across the energy sector, from grid optimization to demand forecasting. However, the infrastructure that powers AI requires significant amounts of electricity. As more industries adopt AI, energy and sustainability leaders must balance growing demand with the need for efficient and lower-carbon operations.

GUEST EDITORIAL

When Utility Fleets Become Operational Risks

Michele Way, Element Fleet Management

For utilities of all sizes, the work required to maintain company vehicles can sneak up on their managers. What starts out as a small collection of work trucks can quickly turn into a diverse fleet of specialized vehicles that support nearly every aspect of field operations and require much greater oversight. For utility managers, it’s crucial to recognize that while fleet sizes grow in size and complexity, so should the processes used to manage them.

POWER POINTS

HOW AI IS LIGHTING UP THE ELECTRIC ENERGY SECTOR

ELISABETH MONAGHAN
Editor-in-Chief

It wasn't long ago that artificial intelligence was more of an abstract yet exciting concept in development but not quite tangible.

Today, AI is no longer abstract. In fact, we're seeing everyone from large corporations to small businesses deploy AI to troubleshoot problems, increase productivity, analyze data, write marketing or creative briefs and serve as customer service chatbots. In the electric energy industry, AI has moved from strategy and possibility into a viable presence embedded in field forms, invoice validation, fault detection systems and customer engagement platforms.

In the past two years alone, we've had several subject matter experts write articles for EEOnline about how organizations are or should be using AI in the electric utility space. In this Q3 issue, several articles discuss AI's role in electric energy. This collection of articles shows how AI has clearly shifted. For this column, I am highlighting four articles that show the different ways AI has taken hold and is lighting up the electric energy sector.

In "**AI Energy Trade-Off: Refining Grids for a Greener Tomorrow,**" April Miller, with ReHack, frames AI's growth in exactly those terms: a trade-off between opportunity and cost. Miller points out that AI is creating new opportunities across the energy sector, from grid optimization to demand forecasting, but the article really centers on the unprecedented demand for AI energy that this same technology is driving. Miller looks at what's behind that surge, from high-speed cabling to advanced cooling, and she covers the policy side too. Right now, the utility sector is still weighing AI and getting ready for it, not yet putting it fully to work.

The second article is "**Managing AI Within Energy Fleets: Why the Hardest Problems Start After Deployment**," by IOTech's Andrew Foster. Foster notes that the conversation has moved past whether AI belongs in grid operations and is now focused on where and how to deploy it, from predictive battery storage simulations to solar fleet maintenance and edge-based voltage optimization.

His argument is about what he calls the deployment-to-operations gap. A model trained on a snapshot of the world will drift as that world changes, and performance drops are usually small at first until they pile up into a real problem, such as operators who've started ignoring false fault alerts. Fixing this, he says, takes four things: keeping data normalized on an ongoing basis, orchestrating the model lifecycle across the whole fleet, watching performance continuously and drawing a hard line that keeps AI in an advisory role, with humans still making the calls. For Foster, this isn't really about the model; it's about whether the operations team is mature enough to handle one that changes underneath them, and that means data science and operations need to know exactly who owns what.

In "**The Human Grid Asset: Scaling Grid Stability through Gamification**," Alex Hill of Sendero Consulting examines AI's role in the behavioral layer of the virtual power plant, rather than the hardware side that Foster covers. Hill's point is that the "human grid asset," the pooled demand flexibility of engaged customers, stays out of reach if utilities only have data to work with, because data alone doesn't drive action without a sustained feedback loop. According to Hill, AI's role is to close that cognitive gap proactively by processing consumption data: predictive analytics let a utility send customers "challenges" ahead of a grid event, instead of scrambling to react once one occurs. Hill builds this around clarity, motivation and feedback: quests, layered incentives and real-time streaks and milestones, with AI as the piece that lets the utility anticipate load and shift behavior ahead of time. This gamification approach is a narrower application of AI than Foster's fleet-scale focus, but it's still a viable, operational use, not just a reporting tool.

The article featured in this issue's Grid Transformation section, "**How Utilities Track Large Capital Programs: Field Data Collection to Real-Time Balanced Scorecard at Scale**," presents another clear indication of the shift from AI as an abstract concept to a solid role in grid modernization. Written by Hari Vasudevan of KYRO AI, along with CenterPoint Energy's Mythili Chaganti, Jerry Caldwell and Chris Harris, this article explains how CenterPoint is executing a $3 billion-plus grid resilience buildout across the Texas coast. CenterPoint Energy aims to cut storm-related outage minutes by nearly a billion for 2.8 million customers by 2029. Its Balanced Scorecard Command Center consolidates six previously separate data systems, such as scheduling, financials, invoicing and field reports, into a single real-time view, now tracking almost 600 projects with 99.56% invoice compliance.

Instead of rolling out a new tool for staff to learn, CenterPoint built AI into the workflows that already fed the KPI platform. It checks timesheets against rate cards and equipment allocations, catching billing discrepancies at invoicing rather than months later. Project managers and coordinators can ask questions in plain language to pull up project performance, and out in the field, an AI agent handles crew questions, provides context, and walks them through their forms and paperwork. On the analytics side, it is watching timesheets, invoices, and project data as they come in, flagging trends and anything that looks off. The Balanced Scorecard Platform is still the system of record through all of this; AI just speeds up the data going in and helps leadership make sense of what's coming out, with people still in charge of the decisions in between.

Each of these articles, along with the other articles in this issue that discuss AI, traces how the utility sector's relationship with AI has shifted from abstract to something tangible. Miller's "AI Energy Trade-Off" treats AI largely as a concept to be weighed and planned around, driven by the demand it creates for energy infrastructure and by the reporting standards emerging alongside it. Foster's article takes the conversation from pilot to fleet-scale operations, where AI must be actively maintained, monitored and governed once it's running on real equipment. "The Human Grid Asset" extends that shift into a new domain, using forecasting analytics to activate the behavioral layer of the virtual power plant. And CenterPoint's Balanced Scorecard system is the most concrete example of that shift: AI is already embedded in timesheets, field forms and invoice validation on a $3 billion active capital program, operating within a platform that remains firmly under human authority. Combined, these articles show a sector that has largely moved past the question of AI's place in utility operations, and is now concentrating on making it work reliably, day to day.

If you would like to contribute an article on an interesting project, please email me:

Elisabeth@ElectricEnergyOnline.com

Elisabeth

THE GRID TRANSFORMATION FORUM

HOW UTILITIES TRACK LARGE CAPITAL PROGRAMS:

FIELD DATA COLLECTION TO REAL-TIME BALANCED SCORECARD AT SCALE

CenterPoint Energy and Its More Than $3.0 Billion Grid Resilience Buildout

HARI VASUDEVAN, MYTHILI CHAGANTI, JERRY CALDWELL AND CHRIS HARRIS

When large-scale utility capital projects and programs arise, the topic of Key Performance Indicator (KPI) monitoring progression management tends to follow closely. And after decades of working on such projects and programs in the industry, the recurring theme is that the data needed to effectively manage them often lags behind the pace of project execution. This leads to contractor and vendor performance issues, invoicing inconsistencies and overbilling, waste, fraud and abuse that end up affecting ratepayers and local communities.

Somewhere between the project management system, spreadsheets, text messages, email threads, phone calls, finance platform and the boardroom, the strategic direction gets muddled and leadership loses the clear and complete picture.

The culprit isn't a lack of technology, though. The missing link is the metrics—what gets measured and why. This is precisely what CenterPoint Energy has set out to solve, and why this provides a solid playbook for utilities on a similar journey.

The structural blind spot in utility capital portfolio oversight

CenterPoint Energy's largest grid resilience infrastructure plan to build the most resilient coastal grid in the country is currently underway, with the Public Utilities Commission of Texas (PUCT) approving the more than $3 billion investment. CenterPoint Energy will continue its investments in the electric transmission system, including elevating more than 99% of substations above the 500-year floodplain and upgrading or rebuilding more than 2,200 transmission structures to improve overall reliability and resiliency.

The purpose is to reduce storm-related outages by almost one billion minutes for 2.8 million customers by 2029. But, with an initiative of this magnitude, how do you keep a pulse on what's happening across hundreds of concurrent projects and programs in real time and in one place?

Most utilities don't. Project managers wait on weekly status reports, which then trickle down to senior leadership in hopefully frequent updates. Schedule delays, missing documents, erroneous figures and all kinds of bottlenecks are painstakingly slow to emerge but quick to throw capital projects and program efforts into chaos. This also creates documentation gaps for future rate case requests.

CenterPoint Energy faces these same challenges. Its resiliency initiative manages a large portfolio of projects and programs across both Resiliency (System Resiliency Plan) and Non-Resiliency (Non-SRP) categories. Multiple contractors, varying contract types (Bid and Time & Equipment) and projects that move through distinct phases, from Engineering and Design, Permitting, to Civil Construction, Substation and Transmission Construction, to Testing and Commissioning and ultimately, In-Service, are normal features.

Meanwhile, project data was scattered across Oracle's Primavera P6 for milestone scheduling, SAP for internal crew financials, Microsoft's SharePoint for stored spreadsheets, Coupa for contractor invoices and bid estimates, email threads and text messages for field updates and contractor timesheets for accrued costs.

Teams were forced to manually assemble all these different pieces just to get the overall view. This is unfeasible when managing hundreds of projects, especially across multiple phases, contractors and fiscal years, with regulators, ratepayers and executive leadership all expecting accountability.

How a decade of field data collection enabled strong program management

In 2016, CenterPoint Energy deployed a field data collection platform to capture what had previously lived in emails, text messages, phone calls and filing cabinets: real-time field updates, field photos, as-built documents, Quality Assurance & Quality Control (QA/QC) inspection reports, daily field reports, contractor performance reports and project progress tracking. For the first time, field crews and project coordinators had a structured way to document what was happening on the ground in real time, rather than reconstructing it after the fact.

That decision created a compounding advantage. Over nearly 10 years of continuous use, the platform has built a depth of operational data that fundamentally changed the relationship between field operations and engineering. The operations-to-engineering-to-standards communication gap that plagues most utility capital programs began to close, not through process mandates or organizational restructuring, but through a shared data layer that both engineering and operations could see and trust.

By the time CenterPoint Energy was ready to build a balanced scorecard for its $3 billion resilience program, it wasn't starting from scratch. CenterPoint Energy built on nearly a decade of structured field data, and that foundation is what made the scorecard possible at the scale and speed the program demanded.

Designing the balanced scorecard framework for program oversight

So, CenterPoint Energy developed a purpose-built balanced scorecard platform that consolidates six data systems into a single operating picture. That platform now manages almost 600 projects with 99.56% invoice compliance and live visibility into schedule, budget, contractor performance and internal quality metrics.

The main takeaway here is the replicable framework behind the platform for any utility managing large capital portfolios and looking to move from reactive project tracking to proactive, data-driven capital portfolio management of projects and programs.

Before building anything, CenterPoint Energy defined the metrics that would actually drive decisions. The goal was to build a management platform, not a reporting tool—one that leadership could use to identify risk, hold contractors accountable and keep a more than $3 billion investment initiative on track.

Every decision has been shaped by this distinction. The Command Center was structured around two primary views: a "Key Performance Measure" (KPM) for enterprise-level capital portfolio projects and program oversight, and a "Project Summary" for individual project drill-through as a repository of major decisions, with reasoning and documents to support future rate cases.

Why metrics design drives program performance

A couple of decades ago, when utilities started interconnecting renewable energy to the transmission grid, capital portfolio project delivery and execution emerged as a significant competitive differentiator.

The Institute of Electrical & Electronics Engineers (IEEE) Power & Energy Society (PES) Technical Activities Coordinating Committee Working Group, in partnership with IEEE Engineering Management, conducted structured reviews of established frameworks, including the Balanced Scorecard, the European Foundation for Quality Management Excellence model and Retail Performance Management Systems. That body of research, published in a handful of peer-reviewed papers and presented at industry conferences, helped establish the intellectual foundation for how utilities could systematically track and improve capital portfolio project delivery and execution.

What was made clear here, and what CenterPoint Energy's program validates in practice, is that sequence is a powerful factor. First, design the metrics. Then you can move on to building the dashboard.

CenterPoint Energy's Balanced Scorecard framework stems from well-established bodies of standards, applied to the specific demands of utility capital portfolio project delivery and execution:

1. **Safety Performance Metrics** based on Occupational Safety & Health Administration (OSHA) standards. It's the regulatory baseline for every utility program's operations, but not often translated into real-time visibility. Tracking both leading and lagging indicators empowers program managers to proactively intervene rather than record and address them after the fact.
2. **Project Schedule and Cost Metrics** are inspired by the Project Management Institute (PMI) best practices, specifically Earned Value Management (EVM) and the Cost Performance Index. EVM adoption is still fairly low across utilities, where project accounting tends to trail on-the-ground realities by weeks.
3. **Quality & Productivity Metrics** draw from Lean Six Sigma Principles and the International Organization for Standardization ISO 9001 framework, so disciplines are focused on process consistency and defect reduction. In a contract-heavy environment across hundreds of projects, proactive quality management helps distinguish strong performers from those with snowballing issues that only belatedly come to light.

These metrics have enabled CenterPoint Energy to close pressing operational gaps, eliminating delays and opaque project management in exchange for scalable success and maximum visibility. Critically, each of these metric categories depends on reliable field data flowing into the system consistently, the very capability CenterPoint Energy had been building since 2016.

Where AI adds value in utility capital program execution

A Command Center built on integrated data creates a natural foundation for AI.

CenterPoint Energy's -leveraged agentic architecture—AI systems designed around actual utility workflows rather than forcing workflows to adapt to the technology. This matters because most AI implementations fail in production when they don't understand how an organization actually operates. This context engineering approach, structuring how AI understands company rules, project history and operational workflows, is what enables the agents to function as additional resources for project managers and coordinators rather than requiring them to "train" the AI or adapt their processes.

One of the most practically important decisions CenterPoint Energy made was to integrate existing systems rather than replace them, a deliberately practical approach. The team embedded AI capabilities directly into the operational workflows that feed the KPI platform. The result is a set of tools that work alongside field crews and project managers, not in place of them.

On the financial side, AI-driven validation cross-checks submitted timesheets against rate cards and equipment allocations when invoices are submitted, catching billing discrepancies. Project managers, supervisors and project coordinators can use natural language queries to dive deep into specific project performance, using the AI agent as an additional resource that is always available.

For field operations, an AI agent embedded in the daily workflow answers operational questions while communicating relevant project context. It also guides field teams through form completion, task creation and document submission. This matters for data quality as much as efficiency, because structured, consistent data entry at the field level is what makes the command center's metrics reliable at the enterprise level.

On the analytics side, the AI agent identifies trends, anomalies and risk signals across timesheets, invoices and project data in real time. When the KPI scoreboard shows a contractor's utilization rate drifting below target or a cluster of projects showing financial variance, AI agents can help leadership understand why and whether the pattern is emerging elsewhere in the portfolio before it shows up in the aggregate metrics.

The important point is that none of this replaces the Balanced Scorecard Platform's core function as the framework. The dashboards remain the authoritative view. AI operates upstream, improving the quality, speed and consistency of the data flowing into that view. Downstream, AI helps leadership interpret patterns and act on them faster. Human decisions fill the spaces in between.

Pilot, iterate, then deploy: A phased approach to software platform adoption

Instead of attempting to deploy the management system across the full portfolio from day one, CenterPoint Energy selected two pilot projects to validate the approach before full rollout.

The pilot served multiple purposes. First, it confirmed that data from each source system aligned correctly and that milestone logic, financial calculations and contractor metrics were producing accurate results. It also gave the end users, from program leadership to project managers, field coordinators and supervisors, a chance to interact with the dashboards and provide feedback before the system was populated with hundreds of projects.

Essentially, CenterPoint Energy built confidence in the platform before scaling it by piloting the system, gathering user feedback and iterating before expanding to the full capital project and program portfolio. It's not a complicated process, but it breaks away from the beaten path—most utilities surrender to the pressure to prove results fast, pushing most programs to broad rollouts that overwhelm users and expose data quality problems at scale.

Over the course of multiple feedback cycles, CenterPoint Energy refined terminology, updating SRP/Non-SRP labels to "Resiliency/Non-Resiliency Projects." CenterPoint Energy also inverted the variance logic to frame metrics positively, showing projects with variance under 10% rather than over, while aligning financial figures between the KPI and Project Summary dashboards. Dynamic color coding for variance and financial metrics was introduced, and separated change orders from additional funding for cleaner financial reporting.

Every single one of these changes was driven by frontline professionals' feedback, the very teams who would be using the Balanced Scorecard Platform daily. That included project managers, supervisors and coordinators. The result? A platform that reflects how the CenterPoint Energy High Voltage Delivery team actually thinks about project and program management, not how a software vendor assumed they should.

Data challenges utilities will face when scaling KPI platforms

It's worth touching on several challenges that arose during CenterPoint Energy's deployment because these are likely to appear in any large-scale utility rollout.

Historical Data

For the Balanced Scorecard Platform to show meaningful trends, it needed historical project data, not just current-state snapshots. Coordinating with project coordinators to enter 2024 historical data and post-issued for construction drawing records gave the dashboards the longitudinal depth needed for trend analysis, but proved challenging due to the scale and depth.

Multi-year Project Classification

Projects that span multiple fiscal years needed logic to place them in the correct year and phase based on milestone date constraints. This is a classification problem that doesn't exist when you're tracking projects in spreadsheets one at a time and requires a well-coordinated and planned approach to solve.

Canceled Project Noise

Without filtering, canceled projects cluttered the portfolio view and skewed metrics. Adding a filter to show only Active and Completed projects cleaned up the operating picture.

600 Projects, 99.56% Invoice Compliance: What the Balanced Scorecard Reveals

The Command Center is in production, and CenterPoint Energy leadership now has visibility into the capital project portfolio and its performance that was previously impossible to assemble and tap into.

Across their current portfolio, the KPI dashboard tracks nearly 600 projects distributed across five phases: Engineering and Design, Civil & Electrical Active Construction, Construction Complete, In-Service and On-time Delivery Performance (measured against a 93% target). This provides leadership with clear visibility into which phases need operational attention.

In terms of finances, 99.56% of invoices are received and paid within the 90-day threshold from task completion to payment. That's a near-perfect compliance rate on payment discipline across hundreds of projects and thousands of invoices. Construction estimate-to-actuals ratio for bid contracts is 100%, while T&E contracts are about 85%, with a target of 93% or higher. This metric informs how CenterPoint Energy manages different contract types and also provides areas for improvement.

The Project Summary dashboard gives leadership drill-through visibility into individual projects. The summary includes historical date-change tracking that shows how milestones have shifted over time, documented change requests with reasons for delays, work order costs by type, personnel assignments across coordination roles and field issues with severity classifications and photographic documentation. This level of project-level detail, consolidated into a single view, is what makes a reporting tool a management tool.

Finally, contractor performance metrics show utilization rates, change order patterns and safety scores by contractor, metrics that enable data-informed conversations around contractor allocation, performance improvement and future awards. Internal quality metrics track design changes post-issued for the construction drawings stage, material quality and corrective actions, construction oversight cost variance and construction performance metrics. These provide leadership with a clear read on where internal processes are performing within tolerance and where they need attention.

Prior to this, these numbers and readings didn't exist in a single view. They were assembled manually, often in scattered and fragmented formats, if they were assembled at all.

The Balanced Scorecard is a systems engineering framework, specifically designed to ensure comprehensive insight across the entire workflow. Leveraging the framework ensures every stage is monitored and optimized for effectiveness, not only helping CenterPoint Energy identify areas for improvement but also to consistently align strategies and objectives with measurable outcomes throughout the entire project lifecycle.

The $1.4 trillion question: Will utilities build visibility before regulators demand it?

Grid investments are accelerating toward **$1.4 trillion** over the next five years to power the AI boom and regulators and ratepayers demand greater accountability for capital deployment. Utilities can choose: continue assembling portfolio status from fragmented systems, or invest in the real-time visibility that makes billion-dollar programs manageable.

CenterPoint Energy proved it can be done. Now, the question for every other utility is no longer whether this level of program control is possible, but whether they'll build it proactively or wait until the next program review exposes what leadership couldn't see.

Hari Vasudevan, PE, is a serial entrepreneur and engineer at the forefront of AI, utilities, vegetation management, storm response and construction management. As founder and CEO of KYRO AI, he drives transformative advancements that enable field services companies to digitize operations, reduce risk and maximize profits. He is also the founder and ex-CEO of Think Power Solutions and serves as vice chair and strategic adviser at the Edison Electric Institute.

Mythili Chaganti leads the execution of CenterPoint Energy's high-voltage grid hardening and modernization initiatives, focusing on enhancing resiliency, reliability and affordability. With over 20 years in the power energy industry, Chaganti also advances sustainable infrastructure and clean energy integration. She is a licensed professional engineer, IEEE PES executive and advocate for innovation in high-voltage delivery.

Chris Harris is a seasoned construction project specialist bringing over four decades of experience in electrical substation projects. Known for turning the seemingly unbuildable into successful realities, combining planning, execution and cost management to consistently deliver projects on time and within budget. Harris is passionate about refining policies, procedures and budgeting practices to elevate performance and create lasting impact.

Jerry Caldwell is a substation resilience construction manager at CenterPoint Energy, with a passion for building resilient, reliable electric infrastructure. He leads capital projects focused on safety, collaboration and operational excellence, helping strengthen the power grid while supporting the communities it serves.

GREEN OVATIONS

THE HUMAN GRID ASSET: SCALING GRID STABILITY THROUGH GAMIFICATION

ALEX HILL

The modern utility landscape is defined by rapidly evolving pressures, ranging from volatile grid reliability and skyrocketing demand to the rigorous implementation of new environmental regulations. One of the most complex challenges is the unprecedented pressure to balance aggressive grid-stability requirements with the absolute need for customer affordability. While billions of dollars have been invested in smart hardware – including Advanced Metering Infrastructure (AMI), sensors and real-time data – grid stability remains at risk because the humans at the end of the wire remain largely unengaged. Getting customers engaged in solutions requires a new approach: activating the behavioral layer of the virtual power plant and cultivating the "human grid asset" - a strategy that turns aggregated demand flexibility into a resource that transforms passive customers into active grid participants.

Traditional consumers' demand response programs are often viewed as cold, transactional exchanges that offer no real, sustainable benefit for the user. For example, simply asking a customer to sweat in their living room for a nominal $5 credit fails to provide the necessary motivation to change daily habits. Consequently, these actions fail to scale to the macro-grid effect needed today. To unlock the full value of smart infrastructure, utilities must move beyond viewing behavioral programs as a secondary concern and begin treating human behavior as a formal grid asset, leveraging behavioral science and gamification to move from just informing customers to actively motivating them.

Information fails to drive action

What many utilities are failing to grasp is the fundamental reality that data alone does not equal action. Simply providing a customer with a bar graph of their energy consumption does not generate the emotional impulse required to change ingrained daily habits or incentivize a meaningful shift in behavior. Psychologically, humans struggle to find value in what they cannot see, and because electricity is both invisible and transactional at its core, a significant cognitive gap exists. Without real-time feedback, it remains difficult for the average customer to connect their immediate, daily actions – such as running a dishwasher at 2 p.m. – to the broader health of the grid or their own personal value.

Similar to how having a fitness app doesn't guarantee a perfect diet, the mere presence of an energy portal doesn't guarantee efficiency. The current "missing link" is a sustained, dynamic feedback loop that goes beyond simple information distribution to actively reinforce positive behavior. Without this loop, the extensive investments in smart infrastructure remain an untapped resource.

Gamification as an operational change lever

Real gamification in the utility sector isn't about making energy "fun." Instead, it serves as a change management tool designed to make optimal grid actions clear, motivating and repeatable. By moving beyond traditional or surface-level incentives, utilities can leverage three primary pillars of engagement: clarity, motivation and feedback. These three pillars, each in their own distinct way, bridge the gap between technical infrastructure and human action.

Transforming customer engagement into a reliable grid resource requires a strategic shift from static data to a behavioral framework built on clarity, stacked motivation, and dynamic feedback loops.

The first pillar, clarity, centers around stripping away the often-confusing technical jargon that has historically defined utility communications and instead providing intuitive "quests" that hold tangible meaning for customers. For example, an effective implementation of an "Off-Peak Challenge" simplifies the objective by replacing complex data with a clear, goal-oriented mission that features a defined start and end point. By framing grid-stabilizing actions in such a transparent way, these tasks become significantly more accessible to the average customer, allowing them to identify exactly what is expected of them in real time. This level of clarity ultimately removes one of the primary barriers to participation in most energy programs.

The second pillar, motivation, focuses on implementing varied incentives to trigger different psychological profiles: economic, social and identity. While many programs rely on incentives associated with only one of these profiles, true behavioral scaling requires a multi-layered approach that addresses the diverse psychological facets that meaningfully drive human action. By stacking these incentives, utilities can create a more resilient engagement model. As the initial novelty of one incentive begins to fade, the others are positioned to pick up and maintain momentum.

- Economic incentives, traditionally used by utilities and expected by consumers, provide immediate financial benefits. Much like a local grocery store offering fuel points to turn a routine grocery run into a game to lower gas prices, utilities can use tangible rewards to transform energy usage into an engaging experience. While **research suggests** that economic incentives can be impactful in the short term for simple, well-defined behavioral goals, they should be paired with other psychological levers to prevent the program from being viewed as merely a transactional exchange. To achieve the macro-grid effect needed today, these rewards must be supported by social and identity-driven incentives to foster long-term, sustainable behavioral change.
- Social incentives rely entirely on community norms and comparative influence to drive action. Humans are inherently social and we have a natural inclination to see how we stack up against our peers, so even invisible energy consumption can become a point of engagement when framed through a social lens. For example, if you have a black belt in martial arts, it signifies your status, changing how people may perceive you and, more importantly, how you behave. By leveraging a similar psychological drive in energy consumption, utilities can unlock a key motivator.

- Identity incentives align with personal values, sustainability and civic contribution. This profile is arguably the most powerful incentive because it taps into intrinsic motivation, which is the internal drive to act in a way that is consistent with how one perceives themselves. For example, a traveler may take an inconvenient route, possibly adding an additional stop on their trip, just to maintain "Platinum Status" with an airline. Additionally, by framing grid-stabilizing actions as a form of civic duty or environmental stewardship, utilities can ensure that positive behaviors are sustained long after the initial novelty of a financial reward or social comparison has faded.

The third and final pillar is feedback, established through real-time loops that mirror the "streaks" and milestones found in successful healthcare wearables. A **2023 study** found that 30-35% of Americans regularly utilize a wearable smart device to track health metrics ranging from sleep patterns to calories burned. Achievement features like the Apple Watch "rings" are particularly effective because they motivate users to reach their goals through quantifiable, real-time measurables. The power of these milestones lies in their ability to transform abstract goals into a series of emotionally resonant events. By mirroring this experience, utilities can bridge the "cognitive gap" with a sustained feedback loop, reinforcing positive behavior and connecting a customer's immediate actions to grid health.

To begin implementing these feedback loops, utility leaders must design digital interfaces that move away from static monthly summaries and toward dynamic, milestone-driven tracking. Operationally, this requires translating raw consumption data into immediate, visual progress indicators, such as localized grid-health "streaks," community milestones or digital badges, that reward ongoing participation. This feedback pattern can provide the essential emotional impulse necessary to turn a general user into a permanent steward of the grid.

Integrating data for real-time action

To drive real-time gamification, utilities must bridge the gap between back-office IT and Operational Technology (OT) to ensure usage data is reflected in the gamified interface instantly. Without this synchronization, the real-time feedback loops required to mirror the "streaks" and milestones of modern digital experiences are impossible to maintain. Alternatively, when unifying these systems, utilities can foster actionable dialogue with the customer.

For data integrity and scalability, utilities should establish robust data governance and cloud-based platforms capable of handling the influx of feedback data from millions of residential endpoints. A centralized, cloud-based approach ensures the platform remains resilient and capable, even during periods of high-velocity data processing, without compromising security or accuracy. Ultimately, this infrastructure is crucial for maintaining a transparent and trusted relationship between the utility and the engaged customer.

Utilities should also implement Artificial Intelligence (AI) and Machine Learning to process complex datasets, allowing the organization to move from a reactive to a proactive state. By leveraging tools like predictive analytics, utilities can be equipped with the capability to issue "challenges" to customers proactively, well before a grid event occurs. Rather than reacting to a crisis, these AI tools allow the utility to anticipate load requirements and incentivize behavioral shifts in advance, effectively turning the customer base into a flexible asset for grid stability.

Navigating implementation – leadership and cultural alignment

Another major component of successful gamification is the dedicated buy-in from utility executive leadership. Leadership must shift the organizational mindset to view behavioral programs not as "marketing fluff," but as a core operational necessity. Like with the implementation of any transformative digital strategy, there is significant risk of change fatigue across the workforce if the vision is not clearly articulated from the top down. Without this cultural alignment, a significant barrier is created, preventing this aggregated demand flexibility from evolving into a truly resilient and scalable pillar of grid stability.

To effectively manage change fatigue, utilities should ensure that digital integration and new customer-facing tools are rolled out with a focus on seamless collaboration between technical teams and customer experience units. This unified approach to cross-departmental collaboration breaks down the traditional silos that often hinder the deployment of strategic initiatives. Utility leadership must create and ensure a streamlined implementation process to prevent the rollout of fragmented tools that overwhelm staff and confuse customers.

Furthermore, as the grid becomes increasingly interconnected, the prioritization of cybersecurity within this engaged ecosystem is paramount. In an active consumer digital landscape, utilities must prioritize protecting customer data and intellectual property by implementing rigorous protocols. This includes adopting a zero-trust architecture to verify every access request within the network. For example, a technician accessing grid data from a mobile device should be routinely verified through identity-based permissions to ensure that a single compromised device cannot threaten the broader system.

Utilities should also utilize end-to-end encryption for all data in transit. This eliminates the risk of exposing sensitive usage patterns – such as data indicating whether a customer is or is not at home – which is essential for maintaining the transparent, trusted relationship necessary for a behavioral program to succeed. Another vital security tool is multi-factor authentication, which adds a critical layer of defense alongside automated threat hunting to proactively scan the network for any indicators that the system could be compromised. These security layers allow utilities to manage the risks that come with increased digital connectivity while ensuring that the behavioral layer of the virtual power plant remains a resilient and secure pillar of grid stability.

Looking ahead

While billions of dollars continue to be invested in smart grid hardware, the transition to a cleaner, more resilient grid cannot be achieved through infrastructure alone. While modern infrastructure generates a wealth of actionable data, that intelligence remains stranded if the end-user stays disconnected from the process, leaving the full value of those extensive investments untapped.

As we look toward an increasingly decentralized energy future and mounting pressure on utility companies to balance decarbonization with affordability, the role of the customer will become more critical. When behavior is intentionally designed and treated as a formal grid asset, utilities can transform their consumer base from passive users into flexible, reliable resources for grid stability. Utilities that master the "how" of behavioral integration will see enhanced resilience, higher customer satisfaction and a sustained ability to meet aggressive environmental goals. Ultimately, the most powerful hack for the modern grid isn't just the implementation of more expensive hardware, but the intentional engagement of human behavior.

Alex Hill joined Sendero Consulting in June 2016 and currently serves as senior manager. In this role, Hill leads project management and system implementation initiatives across the electric utility, healthcare, financial services and oil and gas industries. Notably, Hill recently leveraged her extensive project management expertise to guide the execution of Sendero's largest-ever IT application implementation. Her industry impact is further highlighted by her recent collaboration with a major Texas utility, during which she helped establish the program management disciplines needed to successfully execute a $3.4 billion System Resiliency Plan. Hill holds a Bachelor of Science in accounting & marketing from the University of Oklahoma.

THE POLE PROBLEM: CLOSING THE FIELD GAP IN WILDFIRE MITIGATION BEFORE THE NEXT SEASON

DAVOUD ZAMANI

Ask a utility asset manager to describe their wildfire mitigation program, and the answer arrives in a familiar order. Vegetation management is first, the largest line item and the one regulators ask about most. The second line item is the covered conductor and the third and fourth are fast-trip settings and public safety power shutoff protocols. Then, at the far end of the capital plan, undergrounding: the permanent fix, California's three large investor-owned utilities **price between $1.85 million and $6.07 million per mile**.

Somewhere in that sequence, often as a footnote in the inspection budget, sits the wood pole.

Undergrounding works. It is also, by the numbers, a generational program. PG&E launched the country's largest such effort in 2021, targeting 10,000 miles. It reached 1,000 completed miles in October 2025, and projects 1,600 by the end of 2026, an estimated 18% reduction in system wildfire risk. That is four years of the most aggressive undergrounding push any U.S. utility has attempted. It is also 16 of the goal.

This ordering is understandable. It is also increasingly hard to defend on risk terms, and regulators have begun saying so. In the 2021 Brewer Fire in Grass Valley, the CPUC's Safety and Enforcement Division **cited PG&E for six violations**, among them that the utility found woodpecker damage exceeding its own replacement threshold on the ignition-point pole and patched a single hole instead of replacing the structure, then let a 12-month replacement work order lapse without completion or reassessment. The fire burned 5.5 acres of Tier 2 high fire-threat district. On a windier day, the same pole and the same deferred work order would have produced a different headline.

An estimated 160 to 180 million wood poles are in service across the United States, and a 2026 review of utility pole failure in **Engineering Failure Analysis** notes that roughly 80% of North American utility poles are wood. Much of that fleet is operating deep into, or past, the 35-to-50-year service life that treated wood is credited with. Their structural condition is not a maintenance detail. It is a measurable wildfire variable that belongs in the asset risk register alongside conductor type and vegetation clearance.

The pole is both a fire cause and a fire casualty

The prevailing mental model treats the pole as a fire victim: something a wildfire happens to, where the question is whether the structure survives flame contact and the countermeasure is fire resistance. That model is not wrong. It is half the picture, and the half that arrives second.

A pole fails in two ways, and they compound. The ignition mode is mechanical. Wind loads the pole, the pole breaks, the energized conductor falls into cured grass and ignition follows in seconds. This requires no fire beforehand, only that residual strength has fallen below the load the wind imposes. The propagation mode is thermal. Once a fire is burning, whether ignited here or arriving from elsewhere, flame contact consumes poles along the corridor, spreading the fire and turning one ignition into a multi-day rebuild across miles of line.

The modes share a root cause and a location. A pole with a decayed groundline is both more likely to break under wind and less able to survive flame contact, because the same lost cross-section that reduces bending capacity also reduces the material available to char before collapse. Decay does not choose one mode. It degrades both.

This matters for how utilities evaluate interventions. A treatment that restores capacity but burns addresses ignition and abandons the asset during propagation. A treatment that resists fire but does nothing for capacity fails in the other direction, since the pole never reaches flame contact, having already broken. The requirement applies to the same asset in the same zone.

Residual strength is the variable that decades of deterioration erode:

- **Groundline decay.** The zone from six inches above grade to two feet below is where moisture, oxygen and soil organisms meet. Fungal decay hollows the pole from the inside out, often leaving an intact-looking shell over a compromised core. Bending stress peaks at groundline, making it the worst place to lose material. It is also where surface fuels concentrate and flame contact is most likely.
- **Woodpecker damage.** Cavities remove material and open paths for water intrusion, accelerating decay, as Brewer illustrates.
- **Checking and splitting.** Wet-dry cycling opens lengthwise checks that admit moisture, reduce shear capacity and give flame a path into the pole's interior.
- **Load accumulation.** A pole set decades ago to carry a primary conductor and a neutral now also carries telecom, fiber, transformers and reclosers. Every attachment adds weight and wind-catching area, so the load has grown while decay has quietly removed the wood resisting it. Demand rises, capacity falls and eventually they meet.

The picture from the engineering literature is consistent: a pole retaining a fraction of its original strength can stand for years without incident, then fail in the first wind event exceeding its degraded capacity. Standing is not the same as sound. That gap, between what an asset register records as "in service" and what the asset can withstand on a red flag day, is the field gap this article is about.

Why the inspection cycle is not enough

Inspection cycles are set by state rule or internal standard. California's **General Order 165** is representative: annual patrols in urban and high fire threat areas, detailed inspections of overhead equipment every five years and intrusive inspection of wood poles over 15 years old on a ten-year cycle, stretching to twenty once a pole has passed. Three structural problems limit what that delivers.

First, the interval is long relative to the hazard. A pole presumed sound for two decades accumulates attachments and decay throughout, and the annual patrol filling the gap is a visual check that cannot see groundline section loss at all. Decay is not linear; once a fungal colony is established and moisture retained, section loss accelerates.

Second, the queue is longer than the season. Brewer was not a detection failure. PG&E found the damage, wrote the work order, and did not complete it. Utility backlogs are structural: after a 2019 incident in which 26 poles failed, **Seattle City Light committed to clearing roughly 6,000 poles** needing replacement or reinforcement at 1,000 to 2,000 annually, constrained by crews, materials, and budget. Southern California Edison **budgeted $1.1 billion over three years** to inspect 1.4 million poles, targeting 35,000 replacements a year. Replacement runs at a few percent of the fleet annually, set by crew capacity rather than capital, and a wildfire-driven acceleration does not fit inside it.

Third, the assessment gives a condition grade, not a risk score. Knowing a pole has lost 30% of its groundline section is useful. Knowing whether that pole, at that span, in that fuel and wind regime, is a candidate for catastrophic ignition is a different question, requiring structural conditions to be layered against wildfire consequence modeling.

Prioritization: intersecting three data layers

The utilities making real progress are not treating pole hardening as a fleetwide program. They are treating it as a targeting problem, aimed at the intersection of three layers most organizations already own but rarely join: structural condition (inspection records, decay measurements, remaining strength, attachment load), usually the most complete layer and the most siloed; ignition consequence (fuel type and continuity, slope, fire perimeters, wind corridor modeling, downwind exposure), typically living in the wildfire mitigation group on another system and refresh cycle; and circuit criticality (whether the span stays energized in high wind, the protection scheme, whether de-energization is viable).

The intersection of all three should be far smaller than any one layer alone: poles simultaneously degraded, sited in high-consequence fuel and on circuits that stay energized in wind are a subset of a subset of a subset. Utilities that have not run this join often assume the hardening problem is fleet-wide, when the acute-risk population may be small enough to address in a single construction season. The fleet-wide framing produces paralysis; the intersection framing produces a work plan. The CPUC has pointed the same way, observing that **with climate-informed modeling down to individual poles**, utilities can locate the most at-risk segments and decide which warrant undergrounding. The same logic applies to hardening, at a fraction of the cost.

Reinforcement as a capital strategy, not a stopgap

Once the target list is defined, the question is what to do with those poles. Replacement is the default answer and the wrong default for most of the list.

Reinforcement, the category of restoring or augmenting a degraded pole's capacity in place, has existed in utility practice for decades, as steel trusses and stubs and remains in routine use: Ohio Edison's inspection program explicitly **reinforces rather than replaces some poles**. What has changed is the range of approaches, the quality of validation data behind them and the recognition that reinforcement can be a permanent asset decision rather than a bridge to replacement.

The industry already has a precedent for this logic in covered conductor. The Public Advocates Office notes that, per project, it takes **one to two years to install against three to four years for undergrounding, at roughly a third of the cost**, letting a utility protect nearly four miles for the price of burying one. The CPUC has called **aggressive covered conductor installation its preferred strategy** on the grounds that it delivers ignition risk reduction comparable to undergrounding at lower cost.

Reinforcement offers the same trade at the structure level. It is typically a same-day, small-crew operation with no customer outage, so throughput per crew-day is a multiple of replacement throughput, at a fraction of the unit cost. Both ratios are site- and method-dependent and belong in a utility-specific business case rather than a rule of thumb. Two advantages matter disproportionately: many of the highest-risk poles sit on slopes, in backcountry right-of-way, at the wildland interface, exactly where replacement is most expensive and a small footprint counts most. And in-place work avoids excavation, pole disposal and easement renegotiation, shortening the permitting path.

Because the pole fails in two modes, any approach worth evaluating has to answer for both, plus the decay driving them.

On **structural capacity:** what is the restored bending capacity relative to the original design class, how is it measured, does it address the groundline zone where stress peaks and how does it compare against the wind loading the span sees with today's attachments?

On **fire performance:** how does the treated pole behave under direct flame contact, against what test standard, does fire exposure degrade the structural claim and does the method protect the groundline zone where surface fuels gather?

On **durability:** does it exclude moisture and arrest decay, or encapsulate a decaying member and defer the problem? What is the reinforcement's service life against the host pole's remaining life? Is there third-party test data at recognized standards, from protocols resembling field conditions?

A method answering well in one group and poorly in another is not a partial solution. It is a solution to one failure mode and an open exposure on the other.

The documentation standard is rising, and it sets the accounting treatment

This is the part of the discussion that has changed most sharply in recent years. Regulators, and increasingly insurers and capital markets, are no longer satisfied with evidence of expenditure. They want evidence of outcome: risk reduction traceable to specific assets, supported by validation that did not originate with the vendor selling the solution. The organization needs to demonstrate, asset by asset, the pre-treatment condition, the intervention applied, the capacity restored, the test data behind the claim, and how that maps to a modeled reduction in ignition probability on that span.

There is a second reason that discipline matters, and for capital planners, it may be decisive. Under **FERC's Uniform System of Accounts**, work extending an asset's service life or capacity is capitalized to Account 364, Poles, Towers and Fixtures, entering rate base and depreciating over the asset's useful life. Repair that does not extend life is expensed to Account 593, Maintenance of Overhead Lines and deducted in the year incurred. The dividing line is not the technique or the dollars spent. It is whether the work is planned, programmatic and life-extending.

That distinction has teeth. SCE's **Pole Loading and Deteriorated Pole programs** are structured, systematically planned, capital-tracked efforts, treated accordingly through the CPUC's rate case process. Brewer sits at the other end: a patch on a pole, internal standards said to replace, work order lapsed, no program wrapped around it. Reactive, non-life-extending, and when it went wrong, indefensible.

So, the same physical intervention lands on either side of the line depending on the program around it. Reinforcement performed ad hoc, as emergency response to a failed inspection, looks like maintenance. The same reinforcement against a prioritized target list, with documented pre- and post-treatment conditions, validated capacity restoration and a defined service-life extension, looks like a capital program. The engineering is identical. The regulatory treatment, and therefore the funding path, is not.

The targeting analysis and the documentation chain are not overhead on the real work. They are what make it fundable, and the same records will satisfy an insurer, a rating agency and eventually a litigator.

The window is the point

Here is the arithmetic that should sit at the center of every Q4 capital conversation. Undergrounding is the durable answer and, at the pace the leading program is achieving, a generational one. Full replacement is capacity-constrained in a way that no budget increase resolves quickly, because the binding constraint is crews and permits, not dollars. Vegetation management is necessary and does not address the structure. Every one is worth doing. None closes the exposure between now and the next wind event in cured fuel.

The pole fleet is the field gap: where a well-targeted, fast-deploying intervention can move the risk needle in a season rather than a decade, and where, designed as a program rather than a series of repairs, the spend can be capitalized rather than expensed.

None of this argues against undergrounding, against replacement where a pole is too far gone to reinforce, or against the vegetation program. It argues for sequencing. Harden the intersection set now. Replace on the normal cycle. Underground the corridors where the consequence math justifies it. Build the documentation chain from day one.

The pole has been the last asset in the wildfire conversation for a long time, because it is unglamorous and because the industry has treated it as infrastructure rather than a risk variable. The physics have never agreed with that ordering, and the regulatory environment is catching up. The remaining question is whether utilities move the pole up the priority list on their own analysis, or wait for an event to move it for them.

Davoud Zamani holds a Ph.D. in chemical engineering and material science from the University of Arizona, and has a background spanning engineering, business development and cross-functional leadership across multiple technology sectors. As co-founder and CEO of GridWrap, he guides the company's strategic direction and its work with utilities to strengthen aging grid infrastructure.

BUILDING SUPPLY-CHAIN RESILIENCE IN THE AGE OF THE TRANSFORMER SHORTAGE

JAMES WHALEY

The quiet machine that decides everything

There was a time when ordering a power transformer was a line item, not a risk factor. You specified the unit, issued the PO and built your project schedule around a delivery date you could trust. That time is gone.

Today, a standard power transformer averages roughly 128 weeks from order to delivery and generator step-up units stretch to about 144 weeks, according to Wood Mackenzie's Q2 2025 survey. For the largest custom units, analysts now quote lead times of up to 4 years. Put plainly: a transformer ordered today may not energize until the decade is nearly out. When the delivery of a single steel box dictates the schedule of an entire interconnection, supply chain stops being a procurement footnote and becomes the critical path.

The shift from afterthought to bottleneck is the defining infrastructure story of this decade. And it has forced the industry to confront a new form of resiliency. For years, utilities have invested heavily in grid resiliency: hardening lines, modernizing substations, deploying microgrids and preparing for extreme weather. But now, supply chain resiliency has emerged as an equally essential discipline. If grid resiliency is about keeping the lights on during a storm, supply chain resiliency is about ensuring the equipment needed to build, expand and modernize the grid actually arrives.

A colleague recently joked that he remembers when sourcing transformers was so easy he barely thought about it. "Everything else occupied my attention," he said. "Who would have thought transformers would become sexy?"

Yet here we are: the quietest piece of equipment on the system has become the defining constraint for data centers, grid modernization programs, electrified transportation, industrial load growth and even routine utility capital work.

A single large power transformer can now take up to four years to deliver, and an entire project can wait on it. Credit: [TebNad / Adobe Stock]

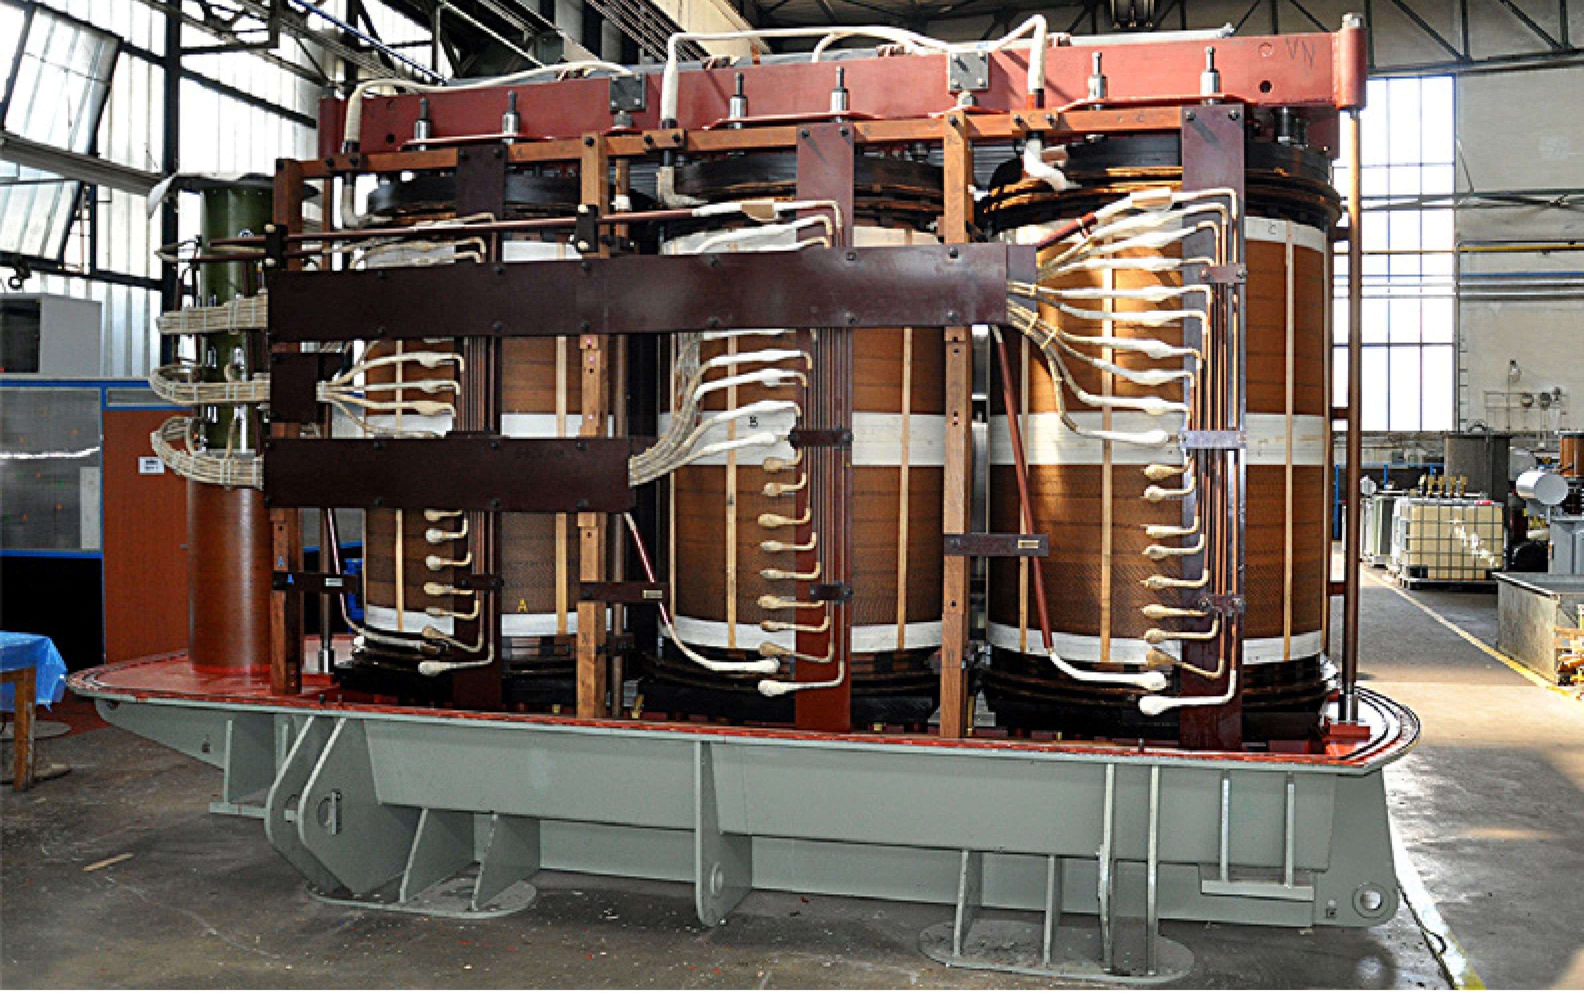

At the heart of every transformer is grain-oriented electrical steel, a specialized material with a single U.S. producer, making it the quiet chokepoint of the entire supply chain. Credit: [petrle / Adobe Stock]

How demand outran supply

It helps to understand the shortage as the collision of three forces rather than a single failure.

1. **Demand exploded all at once.**
   The grid is being asked to do far more than it was a decade ago. AI-driven data centers, the electrification of transportation and industry and aggressive grid modernization programs are all pulling from the same constrained pool of equipment. Since 2019, demand for generator step-up units has grown 274%, and substation power transformers are up 116%. Data centers, in particular, have become one of the most aggressive new buyers, often competing directly with utilities for the same limited manufacturing slots.

2. **The supply base couldn't flex.**
   Transformer manufacturing is a historically low-margin, capital-intensive business and domestic capacity simply wasn't built for this surge. The deeper bottleneck sits at the raw-material level: grain-oriented electrical steel (GOES), the specialized steel used in transformer cores. Cleveland-Cliffs is the only domestic producer, meaning that every U.S. manufacturer relying on domestic steel draws from a single well.

   Compounding this, roughly 80% of large power transformers used in the U.S. are imported, exposing critical infrastructure to global trade dynamics and geopolitical pressures.

3. **Prices followed.**
   Constrained supply against surging demand created exactly the pricing power you'd expect. Power transformer prices have risen 77% since 2019, distribution transformer prices have risen 78–95% and generator step-up transformer prices have risen 45%.

There is a useful contrarian view worth airing, because resilience starts with an honest diagnosis. Some brokers argue the "shortage" is partly self-inflicted; it is a procurement problem more than a manufacturing one. They contend that the bottleneck lies less in factory output and more in the structure of utility and EPC procurement: layers of qualification rules, vendor lists and internal hierarchies that prevent alternative suppliers from reaching decision makers. The counterpoint is equally real: large power transformers are highly customized, heavily regulated assets, and many utilities face legitimate constraints tied to standards, testing protocols, cybersecurity requirements and domestic content rules. Both things are true. That tension — between moving faster and meeting legitimate requirements — is precisely where resilience strategy lives.

Building more capacity takes years

The first instinct is to assume new factories will end the problem. They will help, eventually. Hitachi Energy announced over $1 billion in investment, including a new plant scheduled to come online in 2028, and Siemens increased its commitment to $421 million for a transformer factory in Charlotte, North Carolina. In total, nearly $2 billion has been directed toward expanding North American transformer production, with new capacity projected to come online by 2028.

Here's the caveat: new capacity does not help the project you're building today. Projects executing between now and 2028 face the market as it exists today. Resilience, therefore, is not about waiting for relief. It's about operating intelligently inside a constrained market for several more years.

Strategy 1: Buy slots, not boxes

The single biggest mindset shift is moving from transactional buying to strategic sourcing. In a scarce market, the slot in a manufacturer's production queue is the asset and is secured years ahead.

Long-term supply agreements are becoming a competitive tool. Organizations that lock in delivery slots years in advance gain a scheduling advantage that cannot be bought at spot pricing once a project is ready to execute. Five-year projected plans and multi-year agreements are shifting from rare to routine.

Diversification is equally important. A resilient supply chain mirrors the diversification logic the industry already applies to energy supply. Just as a resilient grid blends solar, wind, storage, gas and transmission upgrades to reduce dependence on any single resource, a resilient supply chain blends domestic and international manufacturers, multiple regions and multiple GOES sources. The principle is the same: no single point of failure should be able to stall the entire system.

Engaging interconnection and planning teams earlier is another critical shift. Projects with flexibility on interconnect timing benefit from talking to utilities sooner, so equipment delivery realities shape the schedule rather than ambush it.

Strategy 2: Standardize and stockpile

One reason transformers are so hard to produce quickly is that they are bespoke. Over 80,000 unique transformer configurations exist, complicating standardization and procurement efforts across the industry. Every custom specification is a custom production run.

Industry-wide standardization is difficult. A newer utility in Southern California has very different operational needs than an older Maine utility, but internal standardization is achievable and powerful. Narrowing your own organization's specification variants reduces complexity, shortens lead times and increases the likelihood that a spare or substitute unit can be sourced when needed.

The cheapest transformer is the one you don't have to buy. Digital monitoring and grid-enhancing technologies stretch the capacity of assets already in the ground. Credit: [JD Studio / Adobe Stock]

Resilience-minded utilities are also rebuilding buffers. Storm stock inventories, spare transformer fleets, mutual assistance networks and shared reserves turn a fleet of individual buyers into a pooled, more shock-absorbent system. Regulators are helping clear the path through "sandbox" programs that accelerate approval of modular substations, mobile transformers and spares, cutting review times from years to months.

Strategy 3: Defer demand with digital

This is where "Digital Transformation & Resiliency" becomes literal. The cheapest, fastest transformer is the one you don't have to buy — because you can squeeze more capacity out of assets you already own.

Grid enhancing technologies (GETs) such as dynamic line rating (DLR), dynamic transformer rating (DTR), advanced power flow control, topology optimization and storage as a transmission asset increase grid flexibility and reduce overall investment costs by improving utilization of existing assets.

Utilities are already demonstrating how digital tools can buy time in a constrained transformer market. PG&E (along with several other large utilities) has deployed strategies such as dynamic transformer rating, dynamic line rating and topology optimization to safely stretch the capabilities of existing assets, effectively increasing usable capacity and deferring capital upgrades when new transformers simply aren't available. These approaches don't eliminate the need for new equipment, but they reduce near-term capital costs and help utilities navigate multiyear transformer lead times.

Reconductoring is the heavy-duty cousin. Replacing old wires with advanced high-capacity conductors can increase line capacity by 25–100%, and, in many cases, offers the lowest-cost and fastest option to address capacity shortfalls. In Texas, American Electric Power upgraded two 120-mile-long 345 kV lines, increasing capacity by 40% without new rights-of-way.

Digital supply chain visibility is another emerging pillar of resiliency. Knowing where every critical unit sits in the production and shipping pipeline turns surprises into forecasts.

Putting it together: The resilience playbook

No single move solves a four-year lead time. Resilience comes from layering tactics so that procurement, engineering **and** planning are pulling in the same direction. For a mixed audience juggling all three, the strategies above reduce to a short, durable playbook.

- **Plan earlier and buy slots, not just boxes.** Treat the production slot as the asset and secure it years ahead through multi-year, multi-vendor agreements. The earlier procurement enters the capital-planning conversation, the more options you keep open.
- **Diversify geographically.** Spread sourcing across regions, manufacturers and material suppliers so that no single trade dispute, plant outage, or single-source dependency can stall an entire project.
- **Standardize internally.** Narrow your own specification variants to shorten lead times and widen the pool of spares and substitutes you can draw from when a unit is delayed.
- **Rebuild buffers.** Expand storm stock, spare-transformer fleets and shared reserves, and use mutual-assistance networks to convert isolated buyers into a pooled, shock-absorbent system.
- **Defer demand with digital tools.** Deploy dynamic ratings, topology optimization and reconductoring to extract more capacity from existing assets, while treating these as complements to physical capacity, not permanent substitutes.
- **Make supply-chain visibility permanent.** Track where every critical unit sits in the production and shipping pipeline so disruptions become forecasts you can plan around rather than surprises that derail a schedule.

Taken together, these moves reframe the supply chain from a back-office function into a strategic discipline;, one that determines whether projects are energized on time.

The bottom line

The transformer shortage is the clearest reminder in a generation that the energy transition runs on physical assets that take time to build. And it has revealed something fundamental: supply chain resiliency is now as essential to grid planning as grid resiliency itself. Hardening lines and modernizing substations will not keep projects on schedule if the transformers required to energize them are still years away. The ability to source, secure, and diversify supply is no longer a procurement function; it is a strategic capability.

Without intervention, extended lead times and elevated costs risk becoming the new normal, potentially slowing data center growth, delaying grid modernization programs and constraining electrification. But "new normal" is a choice, not a fate. The organizations that will keep their projects on schedule through 2028 and beyond are the ones treating supply chain resiliency as a discipline: planning earlier, standardizing specifications, rebuilding buffers and using digital tools to stretch the capabilities of the assets they already have. Utilities such as PG&E have shown that dynamic transformer rating, dynamic line rating and topology optimization can safely increase usable capacity and defer capital upgrades when new transformers simply aren't available — but these tools buy time; they do not replace the need for physical equipment.

Most importantly, supply chain resiliency requires global diversification. The United States simply does not manufacture enough transformers to meet current or future demand, and new domestic capacity will not arrive fast enough to close the gap. Domestic content rules and tariffs may shape procurement, but they do not change the underlying math: the grid cannot be built, modernized, or expanded on domestic supply alone. A resilient supply chain blends domestic manufacturers with trusted international suppliers, ensuring that timing, quality, and cost remain aligned with project realities. In a market where a four-year delay can derail an entire business plan, availability becomes a form of resiliency.

The box in the substation yard may be quiet, but the strategy around it can no longer afford to be. Supply chain resiliency (global, diversified, digitally informed and strategically planned) is now one of the most important tools the industry has to keep the energy transition on track.

James Whaley is a seasoned marketing and creative leader with more than 25 years of experience in design, branding and communications, with a strong focus on the energy industry. His work has supported the growth of solar, construction and energy infrastructure companies by shaping brand identities, driving market engagement and developing strategies that respond to the evolving needs of the energy sector. In his current role as marketing & IT manager at GRID Infrastructure Solutions, he continues to craft innovative movements that connect utilities, developers and EPCs with the solutions that power America's energy future.

References

1. **PG&E** — PG&E Powers Ahead on Breakthrough Grid Innovation with Dynamic Line Rating, Asset Health Monitoring Demonstration (Dec 11, 2025), PR Newswire **https://www.prnewswire.com/news-releases/pge-powers-ahead-on-breakthrough-grid-innovation-with-dynamic-line-rating-asset-health-monitoring-demonstration-302639621.html**
2. **IndustrialSage** — Power Transformer Lead Times Hit Record Highs as U.S. Grid Equipment Shortage Deepens (May 2026) **https://www.industrialsage.com/power-transformer-lead-times-us-grid-shortage/**
3. Congress.gov / CRS — Electricity Distribution Transformers: Supply, Tariffs and Policy Options (Apr 2026) **https://www.congress.gov/crs-product/R48933**
4. Columbia SIPA Center on Global Energy Policy — Unlocking Transmission Efficiency through Grid Enhancement Solutions (Apr 2026) **https://www.energypolicy.columbia.edu/publications/unlocking-transmission-efficiency-through-grid-enhancement-solutions/**
5. U.S. DOE — Energy Department Announces $1.9B Investment in Critical Grid Infrastructure (SPARK program, Mar 2026) **https://www.energy.gov/articles/energy-department-announces-19b-investment-critical-grid-infrastructure-reduce-electricity**

MANAGING AI WITHIN ENERGY FLEETS: WHY THE HARDEST PROBLEMS START AFTER DEPLOYMENT

ANDREW FOSTER

Over the last two years, the conversation about artificial intelligence in the electric utility sector has changed a lot. Instead of asking if AI belongs in grid operations, people now talk about where and how to use it. Battery energy storage systems use predictive simulations at the asset level. Solar fleet operators use machine learning to help decide which maintenance tasks to prioritize. Distribution utilities are trying out edge-based analytics for things like voltage optimization, fault detection and demand response across their growing portfolios of distributed energy resources.

The technology is effective, as pilot projects have shown. However, as utilities move from small-scale tests to AI across many DER sites, new challenges are emerging. These challenges are less about whether the models are correct or powerful enough, and more about a question that is often overlooked: how do you keep AI reliable, trustworthy and manageable once it is running across a distributed fleet in real-world conditions?

This is not a small issue. As utilities with DER portfolios grow and become more complex, how they answer this question will determine whether AI becomes a lasting part of their operations or just an expensive technology that loses value after the first year.

The deployment-to-operations gap

Most AI projects in the utility sector follow a familiar trajectory. A use case is identified. Data is gathered and cleaned. A model is trained, tested against historical data and validated in a controlled environment. It performs well. It's deployed to a site or a handful of sites, and the early results look promising.

What happens next is where things get complicated.

The model was trained on a snapshot of the world as it existed during that training window. But utility environments do not hold still. A battery storage site commissioned with one OEM's equipment receives a firmware update that changes the data format for state-of-charge readings. A solar installation adds a new inverter model with a different communication protocol. A distribution feeder gets reconfigured after a capacity upgrade, altering the load profile that the model learned from. A demand response program changes its dispatch rules, shifting the patterns the model uses to make recommendations.

None of these changes are unusual. In a utility managing a growing DER fleet, they happen continuously. Each one has the potential to diminish model effectiveness in ways that are difficult to detect without deliberate monitoring.

In centralized software setups, operational drift is handled with standard DevOps practices, like version control, continuous integration, automated testing and rollback procedures. But AI at the grid edge is different. Edge nodes are spread out over large areas. Many have spotty or no network connections. Hardware varies by site and firmware versions can change over time. Getting physical access to these sites often requires planning and travel.

The tooling and practices that work in the cloud do not translate directly to this environment. And yet, the models running at the edge are making or influencing decisions that affect real grid operations.

What drift actually looks like in DER operations

It helps to be clear about what model degradation looks like in practice. The effects are usually not dramatic. Instead, small drops in performance add up over time.

Take a battery storage fleet where AI models handle charge and discharge schedules to improve market participation and grid support. When the models were first trained, the fleet had lithium-ion systems from two manufacturers using the same protocol. Six months later, a third manufacturer's system is added, but it reports cell temperature and voltage differently. The model keeps running, but now it gets data that doesn't match what it was trained on. Its recommendations for the new equipment become less accurate, and over time, this can affect how it manages the entire mixed fleet.

Or consider a distribution utility that uses edge-based fault detection throughout a set of DER-interconnected feeders. The model was trained on fault signatures collected during a specific seasonal period. As conditions shift, new load patterns emerge from EV charging infrastructure that was installed after the model was trained. The fault detection model begins generating false positives on feeders with heavy EV charging loads, because it interprets the new load signatures as anomalies. Operators start dismissing the alerts. Confidence in the system erodes, and an actual fault eventually gets missed.

These are not exotic failure modes. They are expected results of deploying learning models into environments that change faster than the models can adapt on their own.

The four disciplines that fleet-scale AI demands

Utilities that have moved past the pilot stage and are running AI across multiple DER sites tend to converge on the same set of operational requirements, whether they arrived at them through planning or through painful experience.

Ongoing data normalization is critical at this point. Making data usable is a continuous process that must keep up with changes in the fleet. When a new asset is added, a protocol changes, or a SCADA integration is updated, the data pipeline must adjust. This means the edge infrastructure should support multiple protocols, use automated tagging to maintain semantic consistency and normalize data near the source before it reaches the model.

This is particularly important in DER situations where assets from different manufacturers, installed at different times, under different service agreements, all need to produce data that is comparable across the fleet. Without that consistency, models built on data from one subset of the fleet will not generalize well to others.

The second is lifecycle orchestration at fleet scale. A model that works today will need to be updated as conditions change. It may need to be retrained on new data. It may need to be replaced entirely by a newer version. When something goes wrong, it needs to be rolled back to a previous known-good state. All of this has to happen across dozens or hundreds of edge nodes, many of which are not permanently connected to a central management system.

Managing this process really cannot be left as an afterthought. Ideally, the system should support staged rollouts, allowing a new model version to be tested on a small number of sites before being deployed more widely. It also needs to be able to cope with sites that have unreliable connections, queueing updates and confirming delivery when a node reconnects. Keeping track of which model version is running on each node is also important.

Utilities that are used to manage firmware updates will recognise some of these requirements, although AI model updates can, in practice, affect system operation in ways that firmware updates typically do not.

Knowing what your models are doing, and what they should not be doing

The next key requirement is ongoing monitoring and observability. Deploying a model without checking its performance over time is a bit like installing a protection relay and then never testing it again. Models need to be monitored on an ongoing basis, with metrics that show not only whether they are running, but also whether their results continue to match expectations.

This requires telemetry designed for distributed environments. Not every edge node will be able to stream performance data in real time. Some will need to batch and forward. Some may need to evaluate model effectiveness locally against predetermined thresholds and generate alerts only when something drifts outside acceptable bounds. The monitoring infrastructure needs to accommodate all these scenarios when providing fleet operators with a consolidated view of model health across the entire DER portfolio.

The fourth requirement is enforcing boundaries. At present, AI in utility operations should only provide advice, with humans retaining control over business-critical decisions. Most people would agree with this in principle, but making it work in practice across a distributed fleet is quite a different challenge.

Enforcing boundaries means laying out clear rules at the platform level for what actions an AI model can suggest, what it can do with human approval and what it cannot do at all. These rules should be the same across the whole fleet, not set differently for each node, so local practices stay in line with company policy. As the industry looks at more autonomous AI, making and enforcing these boundaries at the infrastructure level becomes even more important.

Why the edge is different from the cloud

Some of these problems also exist in cloud-based AI deployments. But the edge introduces constraints that considerably alter the calculus.

Latency requirements are the most obvious differentiator. For DER operations that require sub-second response times, e.g., frequency regulation, voltage support, or rapid dispatch adjustments, processing must occur at or near the asset. Sending data to a cloud environment for inference and waiting for a response introduces delays that are simply incompatible with these use cases.

Operational margins are just as important. Edge environments commonly have limited computing power, memory and storage. They might be in places without reliable network connections and may need to run on their own for long periods. Hardware setups are different at each site, so the edge software must work on anything from small gateways to large servers in substations.

All of this means that the AI lifecycle management infrastructure cannot assume a homogeneous, always-connected, resource-abundant environment. It has to be designed for heterogeneity, intermittent connectivity and constrained resources from the start.

The organizational dimension

It's important to recognize that managing AI across DER fleets is not only a technical challenge. There is also an organizational side that is just as important and often overlooked.

Who owns the AI models once they are deployed? In many utilities, the data science team that built the model hands it off to operations, and the operational technology team is expected to manage it going forward. But OT teams are accustomed to managing deterministic systems. A protection relay behaves the same way every time it encounters the same input. A machine learning model does not. Its behavior is determined by the data it was trained on, and when that data no longer reflects the operating environment, the model's behavior changes in ways that are difficult to predict without specialized monitoring.

Utilities that handle this switch well set up clear ownership between data science and operations teams. They define who is responsible for monitoring model effectiveness, when to retrain, and how to handle problems. Utilities that struggle often have gaps in ownership, so no one notices issues until they become obvious operational problems.

Where is this heading?

The trajectory of AI in DER fleet operations is clear. Models will get more capable. Edge hardware will get more powerful. The use cases will expand from advisory analytics into closer integration with control systems. Some of that is already underway.

But none of that progress will matter much if the utility sector does not build the operational infrastructure to manage AI reliably at scale. The models themselves are only as valuable as the organization's ability to keep them current, monitor their performance and sustain trust in their outputs over time.

This is not simply a technology problem. It is a challenge of operational maturity. Utilities that treat AI governance as a key part of operations, like asset management or cybersecurity, will get lasting value from their investments. Those who see it as a project that ends at deployment will end up with more and more models they cannot trust or verify.

The DER fleet of the future will be more intelligent, more autonomous and more responsive than anything the sector operates today. Getting there doesn't require just better models, but better infrastructure for managing the models themselves. That work is less visible and less exciting than a new algorithm, but it is the foundation on which everything else depends.

Andrew Foster is the chief product officer at **IOTech**, with over 20 years of experience developing IoT and Distributed Real-time and Embedded (DRE) software products. He has held senior roles in product delivery, management and marketing, and frequently speaks at industry conferences on distributed computing, middleware, embedded technologies, and IoT. Foster holds an M.S. in computer-based plant and process control and a Bachelor of Science in digital systems.

WHY HALF OF UTILITY AI INVESTMENTS AREN'T PAYING OFF, AND WHAT CHANGES THAT

Closing the gap takes budgeting discipline and engineering discipline, not a better algorithm

SCOTT FOSTER AND EVARARDO CAMACHO

Collaboratively, we've been working with utility operators for decades, and while we're seeing AI adoption in grid operations accelerate, reported ROI is lagging.

Across industries, 79% of enterprises have adopted AI agents in some form, but only 11% have them running in production — the largest deployment backlog in enterprise tech history — while a separate cross-study analysis finds 95% of generative AI pilots fail to deliver expected returns.

But for utilities specifically, the obstacle is rarely the model. It's that AI tools optimize a single narrow task without ever connecting to the broader operational picture.

Where investment gets decided

The most common pattern behind stalled AI investment is that the pilot is funded as an experiment, not as phase one of a deployment. The budget covers the sandbox, the demo data and a few months of vendor support; success is declared when the model produces promising results.

What rarely gets scoped up front is the work that makes AI operational: integration with OT systems; data pipelines from SCADA and AMI infrastructure whose integration interfaces exist but vary widely in format and capability from one utility's implementation to the next; security review; compliance mapping; and the change management that gets operators to trust the output. That second half usually accounts for most of the cost and the calendar time. The technology didn't fail; the procurement was designed around a demonstration instead of an outcome.

The bigger mistake is treating AI as something you acquire rather than a decision about how the organization operates. An organization that talks about AI as something you buy will buy things. A tool for outage prediction here, one for vegetation management there. Each with its own data requirements and vendor relationship.

Eighteen months later, there's a portfolio of point solutions, each adequate in isolation, none compounding, because the real problem lies in the connections between systems, not in any one tool.

That framing gets locked in at budget approval. If the approved objective is to "demonstrate that AI can predict transformer failures," the pilot is already in trouble — that goal can be met fully while delivering zero operational value. The objective that belongs in that document is narrower: reduce unplanned outages by a defined percentage, measured in real time, feeding into the existing work-order system. No AI spend should be approved without a named operational owner, a real business metric and a defined production pathway in the original task.

Measuring what works

Utilities are good at measuring traditional capital investments – an asset deploys, performs to spec, value accrues on a known schedule. AI doesn't follow that line; it arrives on a curve that starts flat while data pipelines get built and operators learn to trust the output, then compounds as that foundation lowers the cost of the next use case. Judged at month six by capital-investment logic, almost every AI pilot looks like a failure.

The initiatives that do reach the control room share a pattern: operations owned them from day one, not innovation or IT; OT, IT, cybersecurity and compliance sat together from the start instead of being consulted sequentially; and governance was written before deployment, not negotiated after an incident. None of this is technically sophisticated; its organizational design was made in the first thirty days.

Why the integration layer decides the outcome

Those organizational fixes only work if the technical layer underneath can support them. Scott's point about procurement lands squarely on an engineering reality: most utility AI pilots are evaluated against data that was cleaned and curated in a lab environment. Production is a different world. SCADA, AMI and outage management systems were built over decades, on different protocols and different assumptions about who (or what) would be reading their data.

A model that performs well against a static dataset can behave unpredictably once it's wired into a live telemetry stream with gaps, latency and format inconsistencies the pilot never had to account for. Interoperability isn't a connector problem solved with an API; it's a data-quality and data-lineage problem that must be solved before a model is trusted to handle anything that touches grid operations.

The second gap is architectural. Utilities operate under compliance regimes (NERC CIP, among them) designed for fixed control systems, not for adaptive models retrained on new data. Bolting security and compliance review onto a finished pilot is where many otherwise-promising deployments stall, because the architecture wasn't designed to answer the questions a reviewer will ask: What happens if the input feed is compromised? How are model updates authenticated and audited, and what's the fallback if the system loses confidence in its own output? Those answers belong on the first architecture diagram, not retrofitted after a security team finally sees the project — by then, redesigning is often more expensive than starting over.

"Production-ready" means something specific in a control room that it doesn't mean in typical enterprise software: deterministic, auditable behavior under conditions the model hasn't seen before, a defined fallback when confidence drops and a clear boundary between what the system recommends and what a human confirms. Building that boundary requires a clean, governed, interoperable data foundation — because a model can't know the limits of its own knowledge if the data feeding it is inconsistent to begin with.

This is also a sequencing problem. Utilities that try to layer AI onto legacy OT/IT architecture in one motion tend to inherit every unresolved data-quality and access-control issue that the architecture was quietly carrying. The model just makes those gaps visible, usually at the worst possible moment. The utilities seeing durable returns invert that order: they treat integration and governance as the deliverable of phase one, funded and staffed like real infrastructure work, with the AI capability layered on top only once the plumbing already works. It's slower up front, and it's the only version of the sequence that doesn't have to be redone.

Half of utility AI investments aren't paying off for two reasons that compound each other: the budgeting and governance discipline to define what production means, and the technical discipline to build for it from day one.

Neither gap shows up in a pilot demo. Both show up the moment a utility tries to turn a promising model into something a control room can rely on. Closing them doesn't require better algorithms. It requires treating AI as an operating model decision with a named owner, a real metric and an integration architecture designed for production before the first line of code is written.

Scott Foster is CEO of Delta Energy & Communications. He has spent more than two decades in utility telecommunications and smart grid technology, including senior sales leadership roles at Southern California Edison and Cooper Power Systems and holds an MBA from the University of Phoenix.

Everardo Camacho is CTO of Delta Energy & Communications, bringing over 15 years of experience architecting and deploying distributed, internet-connected systems at an industrial scale. In previous roles, he led the development of mission-critical infrastructure for the U.S. Department of Defense and the National Science Foundation.

MATERIALS AS AN ASSET MANAGEMENT VARIABLE IN T&D RELIABILITY

DR. NICOLAS IMLINGER

Silicone composite insulators support high-voltage equipment in an outdoor substation, where hydrophobic, weather-resistant housings contribute to long-term insulation performance. (Source: WACKER)

Transmission and distribution assets must perform for decades while facing moisture, contamination, ultraviolet radiation, temperature cycling and mechanical stress. North American load forecasts add urgency to those decisions. **NERC's 2025 Long-Term Reliability Assessment** identifies data centers and other large commercial and industrial loads as the leading source of projected demand growth over the next decade. Equipment manufacturers and utilities often evaluate these demands through electrical ratings, mechanical loads and service conditions. Material chemistry sits underneath each requirement because the formulation and processing of an insulating elastomer determine how the finished component behaves over time.

Silicone rubber has become an established material for insulators, surge arresters, bushings, cable joints, terminations, connectors and coatings for ceramic and glass insulation. Its value comes from a combination of surface hydrophobicity, dielectric performance, weather resistance, flexibility and processability. The label "silicone" covers a broad family of materials. Polymer structure, reinforcing fillers, functional additives, cure chemistry and manufacturing conditions work together to create the performance profile of a specific material system.

For asset management teams, this chemistry has practical consequences. Material decisions made during component development influence resistance to contamination, the consistency of outdoor performance, installation and handling characteristics, maintenance requirements and service life. This article examines how silicone formulation, processing and application design connect materials science with the long-term performance of T&D components.

A performance profile begins with formulation

Silicone elastomers used in electrical applications generally rely on a polydimethylsiloxane polymer backbone. The silicon-oxygen chain gives the material high molecular flexibility, low surface energy and stability across a broad temperature range. These inherent characteristics provide the foundation for hydrophobicity, elasticity and weather resistance, although the polymer represents only one part of the final compound.

Formulators use reinforcing silica to build mechanical strength and control rheology. Mineral fillers can improve resistance to tracking and erosion under electrical stress. Crosslinking agents establish the cured network, while catalysts, inhibitors and other additives control the cure profile and production window. Pigments and functional additives may support identification, flame performance, electrical conductivity or other application requirements. Each adjustment creates trade-offs among processing behavior, mechanical properties, electrical performance and long-term surface behavior.

Cure chemistry adds another layer of control. Heat-cured solid silicone rubber, liquid silicone rubber, extra-low-viscosity systems and room-temperature-curing materials serve different manufacturing methods and component geometries. A compound developed for injection molding an insulator housing needs a different viscosity and cure profile than a spray-applied coating or a low-pressure material used around a hollow composite core. Cable accessories may require permanently elastic insulating, conductive and semiconductive materials that manufacturers can mold with precise interfaces.

These differences explain why a material specification should describe the required performance and processing route with greater precision than a generic callout for silicone rubber. A material supplier can help an equipment manufacturer connect service conditions with compound architecture, cure behavior and manufacturing equipment. That collaboration reduces the risk of selecting a material that performs well in a laboratory data sheet yet creates instability during mixing, metering, molding, extrusion or curing.

Hydrophobicity under contaminated conditions

Outdoor insulation accumulates salt, dust, industrial residue and other contaminants according to the local environment. When moisture wets a contaminated hydrophilic surface, dissolved ions can increase surface conductivity. Silicone rubber responds differently because its low surface energy encourages water to form droplets instead of spreading into a continuous film. This behavior can interrupt conductive paths across the surface and support resistance to leakage current and flashover under polluted conditions.

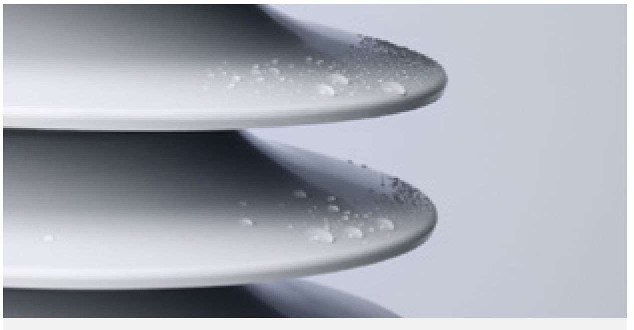

Water beads on a hydrophobic silicone insulator surface, helping limit the formation of continuous conductive moisture films. (Source: WACKER)

Silicone also exhibits hydrophobicity transfer. Low-molecular-weight siloxane species can migrate from the bulk elastomer toward the surface and into an adjacent contamination layer. As that transfer progresses, deposited grime can acquire water-repellent characteristics. This dynamic recovery helps explain why silicone insulation can retain useful surface behavior after contamination or temporary environmental stress.

Formulation controls the rate and extent of that response. The polymer system, filler package, crosslink density and concentration of mobile siloxane species influence hydrophobicity and its recovery after stress. Severe contamination, sustained electrical activity and aging can alter surface behavior, so a credible performance assessment considers both initial hydrophobicity and the material's ability to recover it over time.

This chemistry applies across several T&D component types. Silicone housings on rod and hollow-core insulators use hydrophobic surfaces to support outdoor insulation performance. The same formulation principles apply to silicone housings used for surge arresters and bushings, where hydrophobicity, weather resistance, dielectric performance and process consistency contribute to long-term environmental protection. In every case, the material operates within a complete component whose geometry, interfaces and manufacturing quality also influence performance.

Weather resistance and permanent elasticity

T&D components encounter long periods of sunlight, ozone, heat, cold and repeated temperature changes. Silicone elastomers can maintain flexibility across a broad temperature range and resist degradation from ultraviolet radiation and ozone. These properties come from the chemistry of the siloxane backbone and from a formulation designed for the expected exposure. The resulting material can continue to provide electrical insulation and environmental protection without becoming brittle under ordinary outdoor aging conditions.

Flexibility also supports mechanical durability. Silicone housings can tolerate impact and handling without the brittle fracture behavior associated with ceramic materials. Composite insulator designs can reduce finished-component weight substantially compared with porcelain designs, depending on the construction and duty. Lower weight and a resilient housing can support transportation and installation, especially where access or lifting conditions create constraints. The component manufacturer still determines the final design, mechanical rating and installation requirements.

Permanent elasticity plays a different role in cable accessories. Joints and terminations must maintain intimate contact with cable insulation as load cycles produce thermal expansion and contraction. A properly formulated silicone elastomer can retain radial pressure over a wide temperature range while preserving dielectric performance. Manufacturers can also use conductive and semiconductive silicone grades within molded accessory designs to support electrical field control at material interfaces.

These applications show why long-term performance requires more than resistance to one environmental factor. An elastomer must preserve its electrical, mechanical and surface properties as those stresses interact. Formulation development therefore balances dielectric strength, tracking and erosion resistance, tear strength, elongation, compression behavior, thermal stability and weathering performance for the intended component and processing method.

Processing is part of material performance

A silicone compound reaches the grid only after an equipment manufacturer converts it into a finished component. Processing influences dispersion, cure, adhesion, dimensions and surface quality, which means that production conditions become part of the material-performance equation. The same base chemistry can produce different results when mixing, metering, temperature control or cure conditions vary.

High-consistency silicone rubber supports processes such as extrusion and injection or compression molding. Liquid silicone rubber can provide automated mixing and metering, rapid cure and detailed molding for complex geometries. Extra-low-viscosity systems can fill large or intricate molds at low pressure, an advantage for hollow-core insulator production where manufacturers need controlled flow around a composite structure. Room-temperature-curing materials support coatings and selected cable-accessory manufacturing routes when ambient cure better fits the application.

Rheology influences whether a material fills a mold without trapping air, maintains a coating thickness on vertical surfaces or extrudes with consistent dimensions. Cure kinetics affect cycle time, adhesion development and the risk of undercured or overcured regions. Filler dispersion influences mechanical and electrical consistency, while surface preparation can determine the adhesion of a coating to ceramic, glass, metal or polymeric substrates. These variables fall directly within the shared work of material chemists, process engineers and component manufacturers.

A chemical manufacturer contributes by developing compounds for defined processing windows, characterizing material behavior and helping customers translate laboratory properties into stable production. That support may include rheological analysis, cure studies, adhesion evaluation, processing trials and troubleshooting around mixing or molding conditions. The equipment manufacturer retains responsibility for component design and qualification, while the material supplier provides expertise in the chemistry and process behavior of the elastomer.

This division of expertise strengthens reliability because it keeps each technical question with the party best equipped to answer it. Utilities and asset owners can benefit from the result without expecting a material supplier to design the full component or prescribe operating practices. The most dependable outcome comes from a chain of validated decisions that begins with formulation and continues through manufacturing, qualification and field use.

Coatings as a materials-based retrofit option

Existing ceramic and glass insulators can present a recurring contamination challenge in coastal, desert and industrial environments. A silicone high-voltage insulator coating changes the chemistry of the exposed surface while preserving the underlying insulator. The cured coating supplies hydrophobicity and hydrophobicity transfer, which can reduce leakage current and support resistance to pollution-related flashover.

The coating itself usually combines a polydimethylsiloxane polymer, crosslinking chemistry, selected fillers and additives in a sprayable system. One-component formulations can cure through exposure to atmospheric moisture. The chemistry must deliver a workable viscosity, adequate open time, adhesion to the prepared substrate and a cured film that withstands electrical and environmental stress.

A technician applies a silicone coating to high-voltage insulators. The coating creates a hydrophobic surface that can help reduce conductive tracking and flashover risk in wet or contaminated environments. (Source: WACKER)

Application quality has a direct relationship with coating performance. The substrate must be clean and dry, the material must remain homogeneous during use, and the applicator must control thickness and cure conditions. Adhesion and coverage require verification after application. These requirements stay within the materials and process domain because contamination at the interface, poor wetting, uneven film build or incomplete cure can prevent the coating from reaching its intended properties.

For an asset owner, the practical value lies in adding a durable hydrophobic surface without replacing a structurally serviceable insulator. The decision to coat, clean, retain or replace equipment belongs to the utility and its engineering partners. The material contribution remains clear: a well-formulated and correctly applied silicone coating can improve surface behavior in polluted environments, reduce cleaning demands and extend maintenance intervals when the application fits the asset condition.

Bringing materials science into lifecycle decisions

Asset management professionals gain useful insight by asking how a component's material system supports the intended environment and service life. A robust specification process can examine the elastomer class, cure system, processing method, electrical properties, mechanical properties, weathering data and resistance to tracking and erosion. For hydrophobic materials, the evaluation should consider recovery after environmental or electrical stress alongside the initial surface measurement. IEC TS 62073 describes contact-angle, surface-tension and spray methods for determining the hydrophobicity of insulator surfaces, with each result representing surface hydrophobicity at the time of measurement.

The questions should stay within each participant's expertise. Utilities define system requirements, service environments and operational priorities. Equipment manufacturers design and qualify the component. Material suppliers formulate elastomers, establish processing guidance and explain how chemistry influences performance. Productive dialogue among those groups can reveal where a generic material callout leaves important variables unresolved.

This approach also helps teams distinguish material limitations from component or process limitations. A field issue may arise from an unsuitable compound, inconsistent cure, poor adhesion, contamination during assembly, damage to the finished component or a design that exposes the material to stresses outside its qualified range. Careful failure analysis should preserve those distinctions because the corrective action depends on the mechanism.

Worker safety and productivity enter the discussion through documented material characteristics. Lightweight, flexible components can support handling and installation. Hydrophobic surfaces and durable outdoor materials can reduce maintenance associated with contamination and weathering. Those benefits remain dependent on the finished equipment and the utility's work practices, yet they demonstrate how a chemical decision made upstream can influence field conditions years later.

Chemistry as a long-term asset variable

Silicone materials support T&D infrastructure because chemists can tailor them for distinct electrical, mechanical, environmental and manufacturing requirements. Hydrophobicity, weather resistance, dielectric behavior and permanent elasticity begin with the polymer, then develop through fillers, additives, crosslinking and process control. The same chemistry can take the form of molded insulator housings, cable-accessory components or field-applied coatings because formulation and processing adapt it to the application.

For asset management, the central lesson concerns traceability. Long-term component behavior should connect back to a defined material system, a controlled manufacturing route and qualification data that reflect the intended environment. When utilities, equipment manufacturers and material suppliers maintain those connections, material selection becomes a disciplined lifecycle decision that carries through procurement and qualification. That perspective keeps the article's focus where a chemical manufacturer can contribute with authority: explaining how molecular design and process execution shape the durability of the components that support the grid.

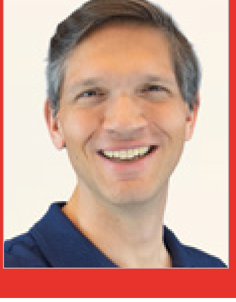

Dr. Nicolas Imlinger is senior director of Energy & Industrial Silicones for North and Central America at Wacker Chemical Corporation, where he leads a regional business focused on silicone technologies for electrification, infrastructure transformation and industrial applications. During more than 20 years with Wacker, he has held leadership roles in Germany and the United States spanning strategic growth, innovation, market development and business operations. Dr. Imlinger earned his doctorate from the Faculty of Chemistry and Mineralogy at Leipzig University.

GUEST EDITORIAL

SEPARATING SIGNAL FROM NOISE:
HOW VERTICAL AI DRIVES HIGH-DEFINITION CUSTOMER ENGAGEMENT

GAUTAM AGGARWAL

For nearly two decades, behavioral energy efficiency programs have revolved around the ubiquitous Home Energy Report (HER) and have become a staple of utility customer engagement.

However, these paper-based, aggregate mailers were built for a different era of technology and customer expectations. Today, vertical artificial intelligence (AI)—purpose-built for the unique data structures and constraints of the energy sector—is helping utilities elevate legacy models to keep modern customers engaged and satisfied.

AI for better segmentation and targeting

Today, vertical AI is helping energy providers modernize legacy platforms to keep modern customers engaged and satisfied. Unlike generic models trained on broad internet text, energy-specific AI is engineered to separate signal from noise in massive streams of smart meter data. It filters out grid fluctuations to isolate the precise energy "signatures" of individual appliances, such as HVAC cycles, electric vehicle (EV) chargers, or water heaters.

This capability unlocks the next level of performance for legacy HERs and energy efficiency outreach. While traditional programs established a baseline for behavioral efficiency, generalized messaging can only take customer action so far. By refining tools like standard neighbor-to-neighbor comparisons (which often paired mismatched households, like a family of five next to a retired couple), vertical AI allows energy providers to evolve beyond demographic-based data and deliver true personalization.

For instance, by continuously mapping usage data to hyper-specific behavioral archetypes in real time, vertical AI can isolate households charging EVs off-peak or flag homes with inefficient HVAC systems whose degraded cycles signal early equipment failure.

So now, instead of sending broad messages to an entire database, energy providers can deliver heat pump rebate offers exclusively to households where they will have the greatest impact.

Building smarter, adaptive customer journeys

Historically, legacy programs from energy providers have relied on one-way communication. A monthly paper report might alert a customer to high energy consumption, but it can lack clear, personal guidance on how to respond, missing the opportunity to prompt enrollment in a demand response program or an optimized EV charging window.

Today, customer communication must adapt to how people actually live. Forward-thinking energy providers are using AI to support two-way engagement and build predictive customer journeys that intercept friction long before it hits a mailbox.

If a home's energy use is pacing significantly higher than normal, the system fires off a proactive high-bill alert mid-month, giving the family a real-time heads-up—and a crucial window to adjust their thermostat—before a budget-breaking bill arrives. Energy providers can pair this alert with a tailored solution based on the home's unique appliance profile. With a single tap, a customer can enroll in a Time-of-Use rate safety net or transition into an affordability program, for example, turning potential anxiety into an empowering experience.

Equipping workforces and educating customers

To bring targeted segments and adaptive journeys to life, customer-facing operational excellence must matter just as much as the underlying algorithm.

For decades, standard energy bills haven't delivered visibility into the factors driving that month's total cost. AI platforms provide that missing transparency, breaking down a monthly total into a clear, appliance-level view of exactly what was spent on refrigeration, laundry, or space heating.

This transparency shifts the dynamic from a lack of information to education, helping customers understand their habits and naturally preventing high-bill phone calls before they happen. When a customer does call, the conversation is faster and more productive because both the caller and the customer service representative (CSR) look at the same itemized data. This enables instant root-cause analysis, allowing a CSR to quickly explain that a specific $40 spike was driven by increased cooling cycles during a mid-month heatwave.

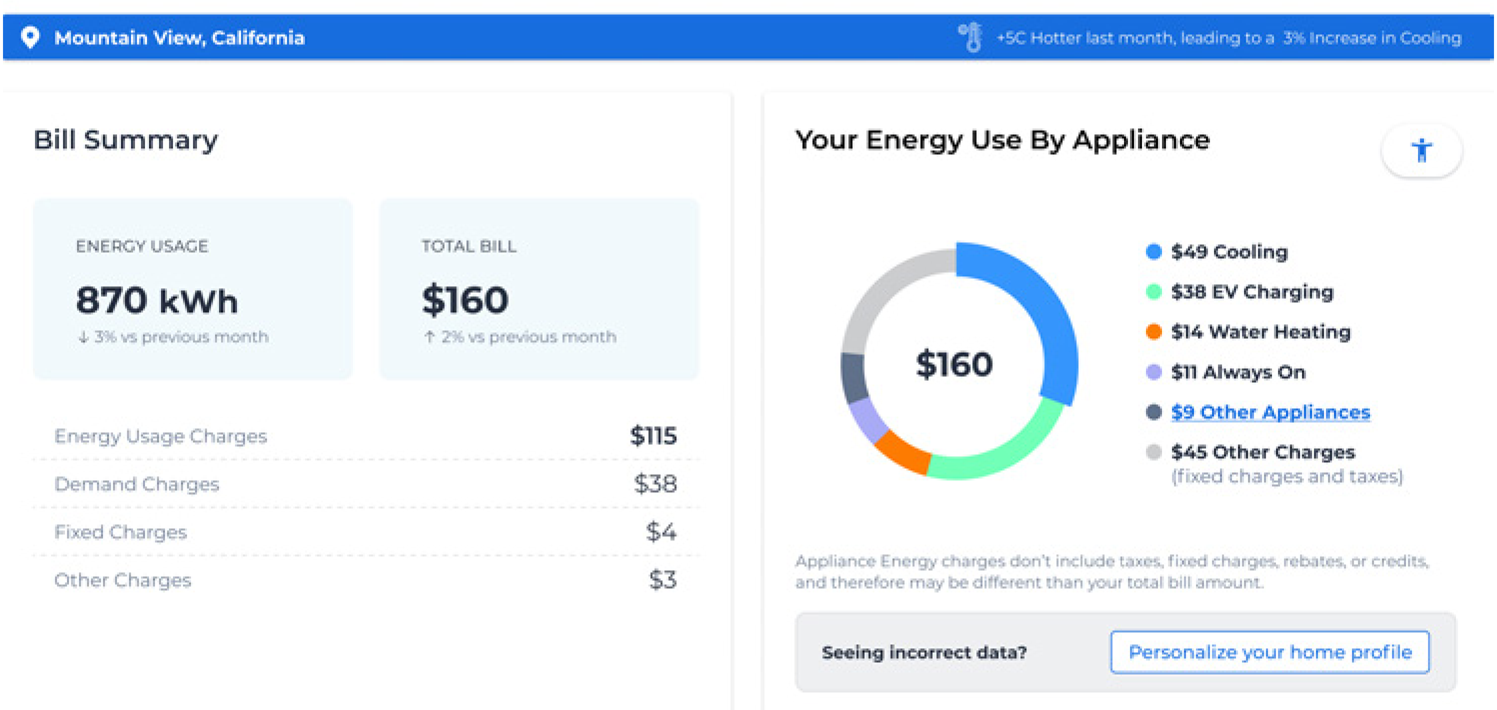

Driving measurable, scalable impact

The evolution from one-size-fits-all outreach to AI-driven segmentation and targeting enables energy providers to build cohesive, multi-layered customer experiences that drive meaningful efficiency in an era of explosive energy demand.

Energy providers will see this come to life across entire service territories, where hyper-personalized touchpoints can, for example, instantly route a customer with poor home insulation to a weatherization assistance program, or offer variable-speed rebates to households running outdated, high-draw pool pumps. As real-time relevance replaces mass-broadcasting, providers are doing much more than sending better reports; they are building a responsive, dual-sided grid where data-driven customer action directly strengthens operational resilience.

Gautam Aggarwal serves as president & Chief Revenue Officer at Bidgely, leading global sales, marketing and revenue strategy. With over 25 years of experience in enterprise technology, he has held leadership roles at FireEye, Cisco and Barracuda Networks. Aggarwal specializes in building high-performing go-to-market teams and driving revenue growth in complex B2B environments.

Gautam Aggarwal serves as president & Chief Revenue Officer at Bidgely, leading global sales, marketing and revenue strategy. With over 25 years of experience in enterprise technology, he has held leadership roles at FireEye, Cisco and Barracuda Networks. Aggarwal specializes in building high-performing go-to-market teams and driving revenue growth in complex B2B environments.

GUEST EDITORIAL

AI ENERGY TRADE-OFF: REFINING GRIDS FOR A GREENER TOMORROW

APRIL MILLER

Artificial intelligence is creating new opportunities across the energy sector, from grid optimization to demand forecasting. However, the infrastructure that powers AI requires significant amounts of electricity. As more industries adopt AI, energy and sustainability leaders must balance growing demand with the need for efficient and lower-carbon operations.

The unprecedented demand for AI energy

AI growth is driving an unprecedented surge in electricity demand. Training and deploying advanced AI models requires massive computational resources, many of which operate continuously in large-scale data centers.

According to research by the International Energy Agency, typical AI-focused data centers **use as much power as 100,000** households. However, the largest ones under construction will likely consume 20 times as much.

This trend presents both opportunities and challenges for the energy sector. AI can help utilities improve forecasting and optimize energy distribution. At the same time, however, the infrastructure powering AI creates significant loads that grids must accommodate.

Energy planners must focus on how quickly demand will grow and how power systems can adapt.

How much energy does AI use?

While individual AI workloads vary widely, the broader infrastructure supporting them consumes enormous amounts of electricity. The annual energy use of U.S. data centers in 2023, excluding cryptocurrency, reached 176 terawatt-hours. This amounted to **4.4% of the country's annual** electricity consumption that year.

Aside from AI training, each query and application deployment requires computational resources. As organizations integrate AI into daily operations, electricity demand will continue to increase.

These realities have intensified discussions about sustainability, prompting governments and tech providers to explore strategies for reducing AI's environmental impact.

Examining the physical infrastructure of data centers

When discussing AI energy consumption, attention often centers on processors and software. However, the physical infrastructure supporting modern data centers plays an equally influential role in determining efficiency and sustainability.

High-speed cabling as the network's nervous system

AI systems rely on the rapid movement of vast amounts of data. Training models and processing workloads require high-performance infrastructure.

This demand extends down to the physical cabling connecting critical systems. Modern networking environments often depend on Category 6a cabling, which supports **data transfer rates of 10 Gbps** at frequencies up to 500 megahertz.

Cabling itself consumes relatively little power, but it forms the foundation of high-speed AI operations. Efficient network infrastructure reduces bottlenecks and improves energy efficiency. As AI workloads become more data-intensive, investments in advanced infrastructure are critical in building efficient data centers.

The critical role of advanced cooling

Cooling is a significant contributor to data center energy consumption.

High-performance processors generate a lot of heat, particularly during AI training and inference operations. Effective thermal management is essential to maintain equipment performance and minimize operational risks.

To address this challenge, computing facilities are adopting innovative technologies, like immersion cooling and intelligent airflow management, that improve efficiency while reducing environmental impact. Some studies indicate that advanced cooling solutions can reduce **carbon emissions by up to 30%** and lower energy footprints by 48%.

The emerging landscape of AI energy policy

Technology is only one aspect of managing AI energy consumption. Policymakers and organizations are increasingly examining how regulation and incentives can support sustainable growth.

The push for mandatory reporting standards

A major obstacle in assessing AI's environmental impact is the lack of standardized reporting.

Many tech companies disclose sustainability metrics, but reporting methodologies can vary. This makes it difficult for regulators and stakeholders to compare organizational performance. As a result, support is growing for mandatory disclosure requirements and standardized measurement frameworks.

The United Nations Environment Programme recommends that countries **establish standardized procedures for measuring** AI's environmental impact and require companies to disclose their operations' environmental consequences.

Incentivizing grid modernization and investment

Electricity grids must support growing computational loads while staying reliable and integrating renewable energy sources. Meeting these demands will require investments in transmission infrastructure, energy storage, smart grid technologies and demand management systems.

Researchers increasingly emphasize the **importance of standardized carbon accounting**, hardware-software co-optimization and life cycle assessments to guide sustainable decision-making.

Exploring pathways to a sustainable AI future

Concerns about AI energy consumption are valid, but significant progress is taking place across the industry, creating opportunities to reduce environmental impacts while supporting growth.

Green data centers demonstrate how operators can improve sustainability in practice. Many governments are introducing policies that encourage the use of renewable energy and lower-carbon infrastructure development.

Among the most notable examples is China's national strategy to **promote green data center operations** over the past decade. Aligning infrastructure growth with environmental objectives allows countries to support digital transformation while managing emissions more effectively.

The success of these initiatives may provide a blueprint for other regions facing similar challenges around AI adoption.

Forging a greener digital world

AI's growing energy footprint presents challenges, but it also creates opportunities to improve efficiency across power systems and digital infrastructure. Through smarter policies and strategic investments, organizations can support innovation while achieving sustainability goals.

April Miller is a Technology Writer at ReHack with more than 5 years of experience covering artificial intelligence and automation.

GUEST EDITORIAL

WHEN UTILITY FLEETS BECOME OPERATIONAL RISKS

MICHELLE WAY

Electric utilities today are facing pressure from every direction. Grid backlogs are causing power projects to wait years to come online, aging infrastructure requires significant attention and increasing severe weather events are placing greater demand on power grids than ever before. At the same time, utilities are navigating supply shortages and skilled worker shortages. There's no question that utility companies are being asked to do more with less.

With so many competing priorities on the table, fleet operations can receive less attention than other areas of the business until problems arise. Yet utility operators cannot afford for their vehicles to slow them down. For utilities of all sizes, from municipals and co-ops to investor-owned utilities, the work required to maintain company vehicles can sneak up on their managers. Especially as their fleet scales in size. What starts out as a small collection of work trucks can quickly turn into a diverse fleet of specialized vehicles that support nearly every aspect of field operations and require much greater oversight. For utility managers, it's crucial to recognize that while fleet sizes grow in size and complexity, so should the processes used to manage them.

Traditionally, fleet management has often been handled through spreadsheets, manual scheduling and decentralized oversight. This approach may work for a short period of time, but as fleets become larger and more specialized, those same approaches can create inefficiencies, increase administrative burden and reduce visibility into one of the utility's most important operational assets.

At a certain point, these DIY management practices will struggle to keep pace with the needs of growing fleets. Building a more resilient fleet operation starts with understanding the unique demands utility fleets face and how complexity can accumulate over time.

The operational requirements of utility fleets

Unlike many other commercial fleets, utility fleets are built to handle essential materials and support specific work tasks, rather than simply moving people or products from one location to another. Every vehicle serves an operational purpose. Bucket trucks are used to access overhead powerlines. Service trucks carry specialized tools and equipment. Even standard vans often function as mobile workspaces stocked with the parts and equipment crews need to do their jobs.

These vehicles also operate under demanding conditions that differ drastically from other commercial fleets. Utility crews often have to travel across large service areas, respond to remote worksites or quickly get to sites with little notice during outages or severe weather threats. During power outages and restoration efforts, when the stakes are at an all-time high, vehicles could be deployed around the clock for days at a time, making reliability and availability absolutely critical to maintaining quality service.

At the same time, utilities have to stay on top of safety standards and regulatory requirements. Routine inspections, preventative maintenance, licensing and compliance documentation all help make sure vehicles remain safe to operate and ready for use in the field. Missing a maintenance appointment or inspection sounds minor in theory, but can lead to unexpected downtime, compliance issues or project delays down the road.

Because of these unique operational demands, every vehicle in a utility fleet plays a key role in keeping crews productive and projects moving. When a utility vehicle is unexpectedly unavailable, the impact of that gap extends beyond transportation. It can delay routine infrastructure upgrades, scheduled maintenance and emergency repairs that customers rely on every day. The condition and availability of utility vehicles can have ripple effects across the entire organization. Understanding this reality is the first step toward building fleet management practices that can grow alongside the company's operations and priorities.

How fleet complexity builds over time

Fleet complexity rarely arrives all at once. It's something that builds gradually as utility operations expand and new demands emerge. A utility that once operated from a single office may, for example, grow to serve multiple districts. Or a fleet of five work trucks can become 15 vehicles with different capabilities, maintenance schedules and reporting requirements. This growth is natural for companies, but as it occurs, the processes used to manage company vehicles must adapt for operations to stay afloat.

It's common to see maintenance records live in spreadsheets, fuel receipts stored in filing cabinets or registration renewals tracked through calendar reminders for growing businesses. But this approach becomes more and more difficult to sustain as operations expand and the company's fleet grows. As the number of vehicles increases, so does the challenge of keeping tabs on fleet activity. Information can become scattered across multiple systems and people, making it more difficult and time consuming to maintain visibility.

This complexity extends far beyond keeping track of vehicle locations. To keep fleets operating efficiently, managers are responsible for scheduling preventative maintenance without disrupting field operations, balancing vehicle utilization across crews, monitoring downtime, planning for asset replacement and understanding the total cost of operating each vehicle over its lifecycle, including tracking weekly fuel costs. The bottom line is ensuring crews have access to the equipment they need when and where it's needed.

There comes a point where fleet complexity reaches a tipping point, causing managers to spend more time managing vehicles than supporting the work those vehicles are meant to accomplish. Left unaddressed, these inefficiencies add up and risk hurting overall business performance.

The hidden costs of unmanaged fleets

One of the biggest challenges of managing a growing utility fleet is many operational inefficiencies aren't immediately visible. A vehicle broken down on the side of the road is an obvious sign that something has gone wrong, but the issues that lead to that moment often develop gradually behind the scenes. Without constant visibility into fleet performance, small inefficiencies can add up over time before anyone realizes their true impact.

Some of the most significant hidden costs include:

- **Uneven vehicle utilization:** Some vehicles may sit idle for days or weeks while others accumulate excessive mileage and wear. Without centralized visibility into utilization, managers may purchase additional vehicles they don't need or shorten the lifespan of others.
- **Maintenance delays:** Missed or deferred maintenance appointments can lead to larger repair bills and unexpected breakdowns down the line.
- **Fuel inefficiencies:** Without consistent monitoring of fleet performance, excessive idling, inefficient routes and unusual fuel spend can easily go unnoticed. This increases operating costs over time.
- **Poor asset replacement timing:** Replacing vehicles too early ties up capital in assets that still have some life in them, while waiting too long increases costs spent on repairs and risks vehicle breakdowns.
- **Compliance risks:** Missed inspections, expired registrations and incomplete maintenance documentation can create safety concerns and expose companies to unnecessary regulatory risk. This can also create a new pathway for unnecessary administrative stress.
- **Increased administrative burden:** Managers that spend a significant amount of time tracking paperwork, coordinating repairs and staying on top of fuel costs for their fleet miss out on supporting their crews in other ways that serve as better uses of their time.

None of these issues drastically disrupts operations all on their own. Together, though, they can compound over time and increase costs, reduce productivity and make it much more difficult for managers and their crews to do their job. All this combined inches toward the tipping point where many utility companies, big and small, realize that managing a fleet isn't simply about maintaining vehicles, but protecting their operation-wide performance.

Recognizing the tipping point

One of the biggest misconceptions about formal fleet management programs is that they are only useful for large organizations with hundreds or even thousands of vehicles. In reality, the tipping point has less to do with fleet size and more to do with fleet complexity. A utility operating ten specialized vehicles across multiple service areas may face many of the same problems as a much larger organization if those assets are still being managed through tedious, manual processes.

The warning signs that a fleet has outgrown existing management practices often appear gradually. Managers may find themselves spending more time catering to their fleet's needs than to any other obligation. Repair costs may increase as preventative maintenance appointments fall through the cracks. Some vehicles may be severely underused while others accumulate excessive mileage and wear. Fuel costs may begin creeping upward without clear explanations. These are all signals that it's time to adopt more structured fleet management processes.

Staying ahead of fleet complexity

The good news is that most of these fleet-related challenges don't pop up overnight, which means utilities have an opportunity to address them before they begin severely affecting operations - and a proactive response is certainly recommended.

Fleet management should evolve alongside the organization itself. The first step is to establish consistent, repeatable processes. Preventive maintenance schedules should be standardized rather than in the hands of individual managers or technicians to remember when a service is due. Vehicle assignments, inspections and compliance documentation should follow clear guidelines that can be replicated across locations, reducing dependence on those spreadsheets and receipts tucked in the back of filing cabinets.

Equally as important is improving visibility into fleet performance. Understanding vehicle utilization, maintenance history, operating costs and asset lifecycles allows managers to make informed decisions about when to rotate vehicles, replace aging assets or adjust maintenance schedules before disruption occurs. Instead of reacting to problems as they arise, with fleet management practices, utilities can identify trends early and allocate resources more effectively. A more proactive approach not only reduces administrative burden but also helps ensure vehicles are available when crews need them.

The goal isn't just to manage vehicles more efficiently, but to strengthen the operations they support. When fleet management becomes a strategic part of utility operations rather than another administrative task, managers are better equipped to continue delivering essential services to the communities they serve.

To get started on revamping fleet management processes, utility companies of all sizes can contact a fleet management provider for customized options.

Michelle Way is the vice president of Element Fleet Essentials at **Element Fleet Management**. She leads Element's small- and medium-fleet division, allowing her to see firsthand the unique challenges growing businesses face and how fleet management can address them.