THIRD EYE ADVISORY™

TECH mahindra

Whitepaper

Power at the Edge

Revolutionizing Smart Grids with Edge Computing

Scale at Speed™

Executive Summary

As global industries pivot toward AI and it permeates every aspect of their operations, it is incumbent on Energy & Utilities (E&U) companies to explore how AI can be used to automate and optimize operations. This includes applying machine learning algorithms for purposes such as building a self-healing energy distribution network or real-time monitoring of distribution and transmission pipelines, as well as upstream equipment in power generation units and oil and gas wells. Near-shore renewable energy generation and the rapid proliferation of EV charging infrastructure have prompted the utilities industry to monitor the LV network. This has increased the number of remote monitoring stations, resulting in a multi-fold increase in data. Edge computing, paired with robust orchestration, offers the low-latency processing, operational resilience, and cybersecurity required for this evolution. This white paper examines the drivers accelerating edge adoption in smart grids, outlines a maturity curve for edge capabilities, and highlights best practices for orchestration and management. Increasingly, industry reports are discussing AI's failure to deliver value and data redundancy at the edge. One way to mitigate this risk is to transition from basic connectivity to AI-enabled intelligence at the edge, utilizing a comprehensive Edge and AI adoption plan that utilities can use to enhance grid resilience, reduce operational costs, meet compliance mandates, and deliver more sustainable, customer-focused energy services.

Introduction

The evolving utility industry landscape is characterized by a growing emphasis on renewable energy sources and adverse impact due to changing weather patterns. This transition is propelled by the widespread adoption of electric vehicles (EVs), increasing energy demands of data centers supporting generative AI and quantum computing technologies, and unpredictable changes in geopolitical policy shifts.
Moreover, in prosumer economies—where consumers both produce and sell energy—this shift in importance highlights the need to transition from centralized, fossil-fuel-based power grids to decentralized micro-grids. These micro-grids integrate traditional power sources with distributed energy resources (DERs), including wind farms, solar panels, heat pumps, batteries, and electric vehicles. This integration of renewables, EVs, and storage systems requires real-time balancing of supply and demand at the local level.

Similarly, the oil and gas industry operates in highly distributed, remote, and data-intensive environments — from offshore rigs to midstream pipelines and downstream refineries.

These advancements have led to a multi-fold increase in the number of sensors in the E&U industry, which in turn generates a substantial volume of operational technology (OT) data. This is transforming the sector by enabling real-time intelligence at the operational frontlines, where traditional network connectivity is limited or costly. According to Gartner , nearly 20% of organizations in the E&U industry have already invested in some form of edge computing, and over 40% plan to do so within the next three years. In this context, it is worth considering whether reliance on cloud solutions, edge computing with orchestration, or a hybrid model is needed.

The answer may not be straightforward. Several factors must be carefully evaluated before making a final decision.

A Short History of Edge Computing

Edge computing emerged as early as the 1980s, initially as client-server computing, and progressed in the 1990s by caching web content closer to users through Content Delivery Networks (CDNs) to minimize latency. In the early 2000s, the proliferation of smartphones and industrial (IIoT) devices exposed limitations in centralized data centers, particularly in terms of latency and bandwidth, while traditional SCADA and PLC systems continued to operate efficiently at the edge. The lack of centralized data using SCADA and PLC was bridged by introducing industrial historians.

The sensitivity of equipment data and the volume of unstructured data have led to a debate over edge-to-cloud versus edge-to-dedicated host. The formalization of edge computing began in the 2010s, as organizations started integrating computing capabilities into equipment described as 'smart'. This 'smart' is in the sensor/actuator that drives the edge, which is different than locally hosted client computing. The extraction of data from these 'smart' devices gradually evolved to include complex computations, resulting in applications in factories, vehicles, and field equipment.

Since 2010, the proliferation of affordable IIoT devices, increased computing power, and advancements in communication networks—including the emergence of 5G—have driven significant expansion within the industry. Major companies have entered the market, leading to substantial investments in this sector. Large to small established players and numerous start-ups alike have introduced advanced technologies such as edge AI, edge analytics and dashboards, edge automation, and autonomous decision-making. The integration of SCADA and PLC systems with sensor data has become increasingly practical, and the ever-evolving (yet never completely solved) IT-OT integration is becoming a prominent focus within the industry.

With this, have we solved the edge problem? No.

Managing hundreds of thousands of edge devices, each running multiple applications and processing petabytes of sensor data, presents significant challenges for industry stakeholders. In the E&U sector, predominantly the utilities industry, which has historically been slower to adopt advanced IT and operational technology (OT) innovations, the pressure to modernize legacy OT infrastructure is mounting. At the same time, organizations must contend with the risks associated with transitioning to complex edge computing environments, including cybersecurity vulnerabilities, integration challenges, and operational disruptions.

What is Driving This Change?

Several factors are driving change in the E&U sector. While CIOs are adopting new technologies and modernizing IT infrastructure, concerns persist over whether existing OT systems can adapt effectively, posing significant challenges. Below are six key reasons the industry is turning to edge computing.

- **Rapid Growth of IIoT and IoT Devices:**
  The surge in these devices has significantly enhanced the capability to monitor a wide range of operating parameters and assess equipment performance. While traditional SCADA/PLC systems continue to track key operational data, storing it alone in a historian may not be sufficient. The increasing volume of information generated by IIoT devices has led to a change in the way this IIoT data gets normalized, analysed, visualized, and used for forecasting. This necessitates robust management and monitoring strategies for both IIoT devices and the time series data that come from these sources. This trend highlights the significance of edge computing, edge applications, and edge orchestration, as local data processing and IIoT device management minimize bandwidth requirements and reduce reliance on cloud infrastructure, thereby providing a more cost-effective and efficient solution.

- **Regulation and Cybersecurity:**
  Since E&U is a critical part of the national infrastructure, it should be safeguarded against cyberattacks, such as ransomware and DDoS. This makes cybersecurity and data sovereignty non-negotiable in this industry. Protecting mission-critical OT systems requires robust, localized measures. Edge solutions support this by segmenting networks, enforcing compliance locally, reducing cloud exposure, securing data handling, and minimizing attack surfaces to mitigate cyber risks.

- **Safety and Environmental Compliance:**
  Worker safety and environmental monitoring require instant detection of pipeline leaks, pressure changes, or unauthorized activity. Edge devices can process video feeds, sensor data, and alarms on-site to trigger immediate alerts — essential for safety-critical scenarios and perimeter security, providing a reliable solution for meeting compliance standards.

- **Operational Efficiency:**
  The increasing global uncertainty is increasing pressure on costs. Improving operational efficiency can help save costs and boost reliability. Edge AI solutions, such as predictive maintenance and remote monitoring, can reduce downtime, extend asset lifecycles, and lower operational and maintenance (O&M) costs. Utilizing AI and ML models trained on industry-specific data enables the analysis of equipment health, the prediction of asset degradation, and the optimization of maintenance expenses, offering a promising future for the industry.

- **Real-Time Decision Making:**
  Edge computing enables immediate data processing and autonomous control, even in disconnected or bandwidth-constrained environments. Downtime in exploration, drilling, or refining is exceptionally costly. Also, in the utilities industry, Traditional cloud-based systems can't meet the latency demands of real-time grid balancing, fault detection, or substation automation.

- **Data Center Cost Pressure:**
  Edge computing alleviates dependence on high-bandwidth, high-cost cloud storage, as well as the recurring expense for running a single LLM in the cloud using a GPU. Data center proliferation has a considerable impact not only on the environment but also on social and economic effects. This has led to widespread criticism, leading to complex geopolitical and government regulations. By processing and filtering data locally, the significant costs and geopolitical pressure associated with transferring large data volumes from the edge to the cloud or centralised data center can be mitigated.

How does Edge Orchestration help the Change?

Understanding Edge Applications

An edge application refers to the functional intelligence deployed close to the data source. It processes raw inputs such as sensor readings, machine telemetry, or substation data into local decisions without relying on a centralized cloud or data center. These applications often run inside containers, virtual machines, or directly on bare metal. They are designed for local data processing and analytics, logic optimization, and low-latency responsiveness, autonomous execution, and efficient resource usage. These edge applications complement or run alongside the traditional PLCs, DCS, and SCADA systems. At the same time, these conventional legacy automation tasks are rule-based, tied to hardware, and human-monitored.

Significant challenges in edge applications include distributing and installing software across hundreds or thousands of edge nodes with varying hardware, operating systems, and network environments, where manual processes can lead to errors.

Additionally, managing updates, rollbacks, and patches for distributed applications requires coordinated version control to mitigate potential downtime or disruption. Many edge environments, such as remote substations and oil rigs, experience intermittent or low-bandwidth connectivity.

Another consideration is obtaining a unified view of application health, performance, and logs across edge nodes, as troubleshooting and diagnostics are difficult without consistent visibility. Consistent security policies, including authentication, encryption, and access control, are necessary across diverse locations. The complexity is further increased when edge solutions are tightly integrated with specific hardware, platforms, or proprietary orchestration tools.

Edge Computing Architecture

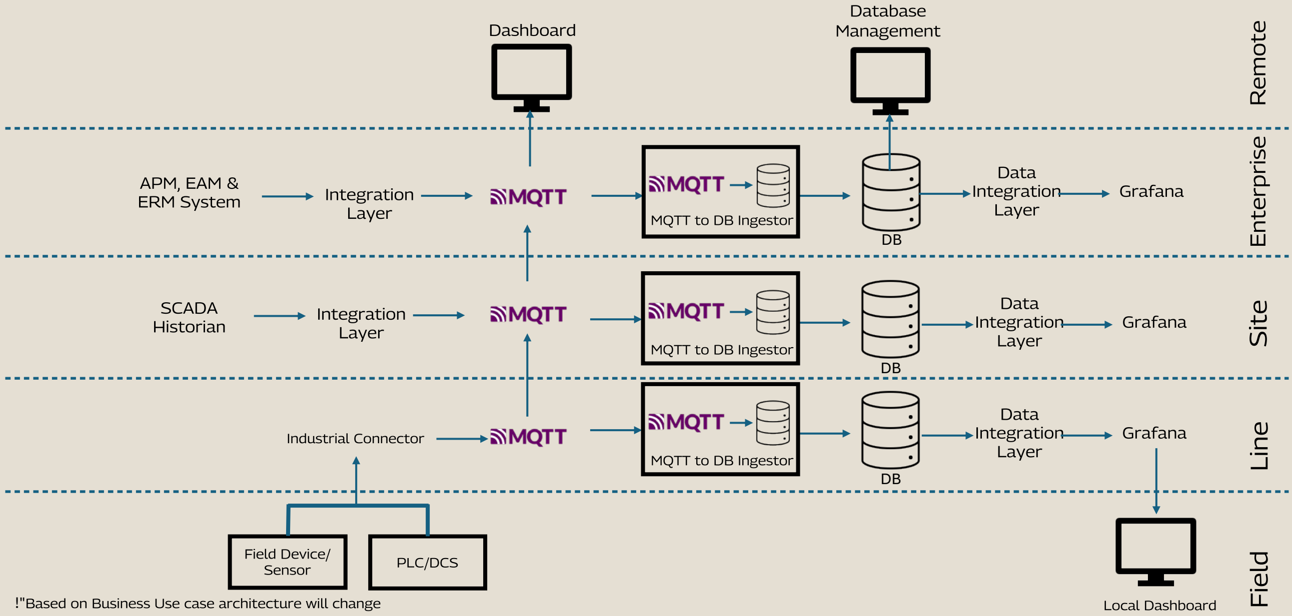

An edge orchestration platform is a management layer on top of the edge application. This platform automates the application deployment, coordination, monitoring, integration, and lifecycle of edge applications across distributed nodes. This can provide consistent security policies across multiple edge devices. It acts as the "control plane" for managing edge computing environments, ensuring that the right applications are running on the right edge nodes securely and efficiently, regardless of location or connectivity.

Industrial edge orchestration platforms, such as Barbara, Schneider Electric EcoStruxure Edge, Siemens Industrial Edge, and Emerson, address issues with centralized deployment, monitoring, and security. They typically use containerization (Docker, Kubernetes) for application isolation and portability, enabling zero-trust security and OTA updates. These platforms are also built for offline resilience and efficient resource usage.

Key Capabilities of Edge Orchestration Platform

- Application deployment and lifecycle management
- Edge device management with unified namespace
- Real-time health status, logs, performance metrics, failure alerts
- Policy-based control and automation
- Edge security and compliance
- Low/no connectivity operations

Opportunities for SIs

Over 40% of E&U industries plan to invest in edge computing within the next three years. Furthermore, the sector is advancing along the edge computing maturity curve, which comprises five widely recognized stages: Independent Edge, Connected Edge, Informed Edge, Reactive Edge, and AI-Enabled Intelligent Edge. Therefore, the challenge is no longer whether to scale, but how to scale it securely, reliably, and across distributed sites. There are significant opportunities for system integrators to assist organizations throughout this transformation journey.

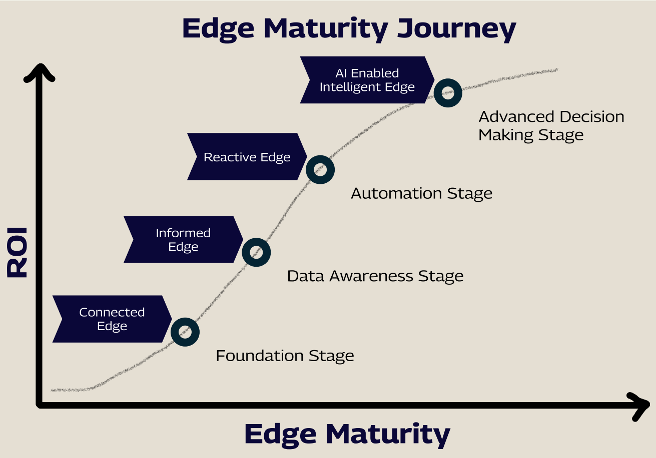

Connected Edge (Foundation Stage):
This stage collects data and alarms from SCADA and PLC systems. Currently, integration of data from sensor devices with existing SCADA data is not yet fully realized. System integrators (SIs), leveraging their expertise in data and analytics, can help industries develop suitable data lakes and analytics engines to visualize and present information to business users.

Informed Edge (Data Awareness Stage):
This stage is crucial for organizations seeking to efficiently manage the massive volumes of data generated by IoT devices, sensors, and other distributed endpoints. By analyzing and filtering data as close to the source as possible, companies can significantly reduce data transmission costs and latency. Pre-processing at the edge allows only relevant, high-value data to reach centralized servers or cloud infrastructure, thereby optimizing storage and computational resources. Effective filtering mechanisms also enhance security by preventing the transmission of unnecessary or sensitive information beyond its intended origin. Challenges persist in standardization, maintaining data integrity, and ensuring robust security across dispersed edge locations—issues that organizations must address as they scale their informed edge solutions. System Integrators can play an important role here.

Reactive Edge (Automation Stage):
It offers an advancement beyond efficient sensor data processing by enabling the configuration of pre-defined rules at the edge to address critical events requiring immediate action. Managing data locally facilitates the development of equipment-specific and process-specific use cases. Edge orchestration platforms not only display critical events through dashboards and warnings but can also issue commands directly to equipment to handle incidents on-site. For instance, self-healing grids leverage this approach to isolate faults, thereby preventing power disruptions.

AI-Enabled Intelligent Edge (Advanced Decision-Making Stage):
This stage takes a step further, leveraging AI to build predictive, generative, and agentic Edges. With a predictive edge, one can create intelligent automation, for example, predictive firm power generation for a combined renewable power generation system, including solar, wind, and a battery bank within the same power plant. By building smaller equipment-specific or process-specific SLM models, one brings in the power of GenAI to the Edge.

Benefits of Advancing Through the Curve

By progressing along the maturity curve, utilities can reduce latency, improve system reliability, lower operational costs, and strengthen compliance. The most advanced stage—AI-enabled intelligent edge—can transform grid operations into proactive, self-optimizing systems that respond dynamically to environmental, market, and asset conditions.

Conclusion

The modernization of power grids demands computing capabilities that operate as close as possible to where data is generated. Edge computing—and consequently, edge orchestration—enables utilities to address the challenges of decentralization, renewables integration, and real-time operational decision-making.

By implementing a structured edge adoption roadmap, utilities can:

- Reduce latency for critical grid operations from seconds to milliseconds
- Improve grid reliability, including reliability of assets in the grid, by enabling faster fault detection and autonomous corrective actions
- Lower operational costs through reduced field visits and automated updates
- Strengthen compliance with cybersecurity and regulatory mandates

The next generation of smart grids will be defined not only by connectivity, but by intelligence at the edge. Utilities that invest now in scalable, AI-ready edge orchestration will be better positioned to deliver resilient, sustainable, and customer-centric energy services for decades to come.

Authors

Premanand Seshadri

Group Practice Head, Energy & Utilities Vertical Practice at Tech Mahindra.

A seasoned utilities industry domain consultant, Prem, enhances industry revenue growth, improves customer engagement, conceptualizes and builds industry-specific solutions, and builds industry-specific partner ecosystems. Prem brings approximately 35 years of utility industry experience, spanning the entire value chain of Power, water, and wastewater businesses.

Krishnakumari Palle

Global Delivery Head, Diversified Industry Groups at Tech Mahindra.

As a seasoned technology leader, Krishnakumari Palle (KK) drives business transformation, IT outsourcing, and multi-service integration for Fortune 500 clients across industries and geographies. She has led large-scale global engagements, overseeing P&L management, high-performance team building, and automation-driven initiatives. KK also serves on the Boards of Directors for Satyam Venture Engineering Services, Mahindra University, and Altavec, actively mentors women in technology, and supports NGOs focused on girls' education. With over 24 years of expertise, she fosters innovation and drives business growth. KK holds an M.S. from BITS Pilani and a B.Tech from NIT Warangal, as well as multiple leadership certifications.

Sahadev Singh

VP of Energy & Utilities Vertical, Europe at Tech Mahindra.

A customer-obsessed P&L leader with a growth mindset and a proven track record for delivering profitable growth across industries and regions. His deep domain experience in E&U is complemented by his passion for building and nurturing high-performance teams. Sahadev brings over 32 years of cross-industry expertise, spanning the E&U, CPG, and Manufacturing verticals.

Dan Miklovic

Expert Advisor -Manufacturing at Third Eye Advisory.

Dan is recognized as a thought leader and one of the top market place influencers in the Asset Performance Management and Manufacturing Operations sectors. His coverage spans across all forms of heavy industry from A&D to Utilities and covering topics such as MES, EAM, QMS, PLM, EH&S, and AI/ML as applied to industrial operations. Dan is involved in both professional and amateur auto sports as an emergency services official and race operations director. He is heavily involved In the start-uc community as a mentor and advisor for entrepreneurship programs at Lindenwood University, Washington University, University of Missouri - St. Louis and the eFactory program at Southwest Missouri University helping clients to validate product strategies and develop business plans to entice investors.

References:

Jaffri, A., Dekate, C., & Choudhary, F. (2023, July 19).
Hype cycle for artificial intelligence, 2023 [Hype Cycle report]. Gartner, Inc.

https://www.gartner.com/en/documents/6352379

About Tech Mahindra

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 152,000+ professionals across 90+ countries helping 1100+ clients, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognises global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit https://www.techmahindra.com/.

TECH mahindra

www.youtube.com/user/techmahindra09
www.facebook.com/techmahindra
www.twitter.com/tech_mahindra
www.linkedin.com/company/tech-mahindra
www.techmahindra.com
mktg@TechMahindra.com