
Cloud-Fog-Edge for Utilities

Successful digital business transformation requires
putting computing power where it's needed

- 04 The Limitations of Cloud
- 04 A Hybrid Approach for Grid Operations
- 05 Utilities Are Pioneering IoT
- 05 Putting Intelligence Where It's Needed
- 06 Fog Computing in Action
- 06 The Case for Edge Intelligence
- 07 Bringing It All Together

The cloud has sucked up a lot of the oxygen in the room when discussing digital business transformation. Businesses have discovered how much money they can save by moving key applications to the cloud, essentially commoditizing traditional IT functions and assets, enabling them to reduce and/or redirect internal IT resources.

The Limitations of Cloud

But the explosion of cloud computing has also exposed some of its limitations, especially those related to the Internet of Things (IoT), which is also growing by leaps and bounds. For businesses that operate extensive networks of connected edge devices to manage critical infrastructure,

cloud computing represents part of the solution, but not all.

More and more utilities are migrating enterprise and back-office functions such as billing, CRM, ERP business applications and even smart metering to the cloud. But the

control and protection of the power grid, with subsecond operational requirements, is another matter. Assuring reliable, efficient and safe power delivery doesn't lend itself to a cloud-only approach.

Cloud, Fog, Edge Definitions

Cloud Computing

The delivery of various scalable, on-demand services over the internet. These services may include software development platforms, servers, storage and application software. The back-end of the services (especially hardware) is managed by a cloud vendor, and a user only pays for services used (memory, processing time and bandwidth, etc.).

Fog Computing

Puts computing and data processing power between the cloud/back office and the edge of a network. This approach can reduce the need for bandwidth by not sending every bit of information over costly or latency-sensitive network circuits, and instead aggregating and analyzing it at certain access points, such as local servers/gateways.

Edge Computing

A distributed computing approach in which computation is largely or completely performed on distributed device nodes or connected edge devices as opposed to primarily taking place in a centralized cloud or network environment.

A Hybrid Approach for Grid Operations

In so many ways, utilities are really at the forefront of IoT adoption because they've been operating de facto IoT networks for so long.

Back when large, centralized power plants provided virtually all the electricity to the grid, delivering electricity to customers was relatively straightforward. Grid operators relied on reasonably accurate demand forecasts, based on weather and historical customer load shapes, to know how much power to generate and deliver. Typically, they have had ample generation resources and reserve margins to draw upon.

But the relative simplicity of matching centralized supply with well understood demand is changing quickly. People and businesses are generating more of their own electricity, creating more supply intermittency, load volatility and power quality challenges on the grid. There are many more grid-connected assets - microgrids, distributed generation, storage, electric vehicles (EVs), controllable loads -

that are capable of generating, consuming, and now storing, electricity. In addition, the number of intelligent, connected devices on the grid that are generating large amounts of data is increasing exponentially. And consumers' expectations for reliable service and new value have never been higher.

The cloud alone cannot solve all these technical and business challenges. A more distributed power grid requires a more distributed IT and computing infrastructure, with data processing capabilities and applications platforms extending all the way to the grid's edge. It also requires evolution from a highly centralized control model, in which operators make decisions from a utility control room, to a more distributed and automated approach.

In this model, artificial intelligence and machine learning enable devices and assets to not only communicate with each other at the edge of the grid, but to quickly take action in response to changing grid conditions to optimize reliability, efficiency,

customer choice and more. Operational challenges such as these require computing and data processing power to be located in multiple levels of the network: in connected devices at the edge of the network (edge computing), at the field-area network level (fog computing) and in the cloud (cloud computing). The optimal location and type of computing power is driven by the operational use cases and what they require in terms of bandwidth, latency, security and other factors.

To be clear, many aspects of IoT are enabled by the cloud. Think of your smart thermostat or smart doorbell with a video security system, both of which you can access remotely with your phone. These are examples of the cloud and edge computing working in concert to deliver new value to consumers. But as more and more devices are added to the network, latency and bandwidth with the cloud become problematic.

Utilities Are Pioneering IoT

How can utilities leverage and utilize this cloud-fog-edge approach to computing power to manage the emerging business and operational requirements of a rapidly changing energy marketplace? The first step for many is to deploy a ubiquitous communications network to support smart metering, distribution automation and other applications such as demand response and distributed energy resource management. This enables the communications network infrastructure to deploy computing power to the fog and the edge. A new generation of network routers, distribution automation equipment, smart inverters and even smart meters now come equipped with

their own computing platforms and data processing power onboard the devices. These advancements drive an array of new capabilities and outcomes that were not possible before:

- The ability to maintain a continually updated and accurate view of which customers are connected to which distribution assets (transformer, feeder, phase) which in turn enables a range of advanced distribution management applications to improve system reliability and efficiency.
- The ability to proactively detect grid hotspots, high impedance connections, downed conductors, broken neutrals and other potential safety problems before they become hazardous issues or costly liabilities.
- The ability to intelligently control and optimize both distributed generation assets and controllable loads at the customer location in near real time to optimize grid reliability and efficiency, protect utility equipment such as transformers from overload, and provide new services to customers.

Putting Intelligence Where It's Needed

Utility operations have always needed back-office computing and now, increasingly, cloud computing to support many enterprise business and customer service functions. From customer billing to call centers to business functions such as load and revenue forecasting, these and many other functions can now be performed most cost-effectively using cloud computing resources.

The growing complexity of grid management in the age of distributed energy resources (rooftop solar, EVs, energy storage) demands that utilities introduce technology to analyze holistically but act locally in real time, often with models developed in the cloud but operated at the edge via artificial intelligence and machine learning techniques. This will help to automate and optimize the low-latency decisions and actions needed to maintain reliability, efficiency, power quality and safety.

Assessing Cloud-Fog-Edge Capabilities

Pros		
Cloud Computing <ul style="list-style-type: none"> • Easy to scale • Low-cost data storage • Enabled by internet-driven global network and robust TCP-IP protocol 	Fog Computing <ul style="list-style-type: none"> • Real-time data analysis • Enables quick action/response • Data remains inside the network • Cost savings on data storage • More scalable than edge computing • Can be operated by IT/OT team 	Edge Computing <ul style="list-style-type: none"> • Puts computing power close to the problem to be solved • Can work without fog and cloud computing available • Local determination of what data should be sent to network or cloud
Cons		
Cloud Computing <ul style="list-style-type: none"> • Higher latency and response time • Higher bandwidth costs • Perceived security concerns • Power consumption • No offline mode • Creates potential data privacy sovereignty issues 	Fog Computing <ul style="list-style-type: none"> • Relies on many links to move critical data from physical assets and devices to digital layer, creating multiple points of potential failure 	Edge Computing <ul style="list-style-type: none"> • Less scalable than fog computing • Often connected through proprietary standards with little interoperability with other systems • Cannot do resource pooling • Operations not extended to IT/OT team

Fog Computing in Action

Outage management represents one area where the introduction of fog computing at the network level can deliver significant improvements in detection and restoration speed. The Edison Electric Institute (EEI) estimates that U.S. smart meter deployment totaled 76 million by the end of 2017, covering 65 percent of U.S. households. EEI projects that figure will reach 90 million by 2020.

A key benefit for the smart meter business case is the fact that each meter can send outage alerts and restoration confirmation signals when power goes out and when it's restored. This capability would help utilities detect, understand and restore power more quickly when outages occur. However, in practice, harvesting that benefit has been more complicated. That's because in any significant outage, the smart meters send a tsunami of outage alerts that can overwhelm

the network and diminish the value of those alerts. By using fog computing at the network level, network routers or localized servers and gateways can filter the large volumes of outage alerts in the field, sending summary information based on an accurate connectivity model that provides accurate and actionable intelligence on outage conditions.

The Case for Edge Intelligence

The introduction of computing and data processing capabilities in edge devices such as smart meters, smart inverters and EV chargers opens up a wide array of opportunities in grid management and new customer services. Edge computing can be used to detect potential transformer overloads at the edge from power flowing in both directions. Corrective action can then be taken locally by reducing controllable loads or summoning/curtailing distributed energy resources at nearby customer sites to reduce transformer loading.

Edge intelligence can be used to activate surgical demand response events based on a customer's ability and desire to contribute. Think of a smart meter that can analyze high-resolution load and voltage data to disaggregate discreet loads within the home, without installing specialized equipment at the customer premises or relying on time-consuming analysis of historical data sets in the back office.

Perhaps the biggest opportunity driven by edge intelligence is the utility's ability to stake a claim in managing local power pools and unlocking the value of transactive energy markets at the edge of the grid. Instead of wasting large amounts of customer-generated electricity during times of low demand and high production, that energy can be bought and sold in localized energy markets. Or it can be intelligently redirected to storage, EV charging or a neighbor who can put it use with a transactive energy app running locally on a blockchain.

Machine Learning and Artificial Intelligence Move to the Edge

We often associate artificial intelligence and machine learning with exotic applications - self-driving cars, speech and facial recognition, robotic control and medical diagnosis - all powered by massive rows of servers full of CPUs or GPUs, at some distant data center. But in fact, AI and ML are getting closer and closer to all of us.

That's because companies such as Google, Microsoft, Nvidia and others have recently introduced technologies and platforms that can cost effectively extend AI and ML capabilities to the edge of the network. Working in concert with cloud services, these devices are capable of processing large volumes of data locally, and enabling highly localized and timely "inference," industry jargon for AI- and ML-driven predictions executed at the edge after having been trained in the cloud; where data storage and processing power are plentiful and scalable.

For example, Coral, Google's new developer platform for local AI, powered by the company's Edge TPU (tensor processing unit) chip, is specifically designed to run machine learning models for edge computing. Coral features a tiny integrated circuit on a credit-card sized development board surrounded by an aluminum housing, with a USB connection to any Linux-based system.

Coral provides a perfect illustration of how cloud, fog and edge work together to put intelligence where it's needed. The heavy lifting for developing a machine learning model and its subsequent training, happens in the cloud where massive amounts of computing and processing power is needed to refine the vast amounts of data required to develop a proper deep learning model.

The result of this training is then compressed and distilled into smaller and faster applications that can be applied quickly to new data at the edge using specialty hardware designed to perform this task. The more data that is collected in the cloud, the better the training models become, resulting in a smarter, faster and more accurate inference - all of which enables edge intelligence to be a true business differentiator.

Bringing It All Together

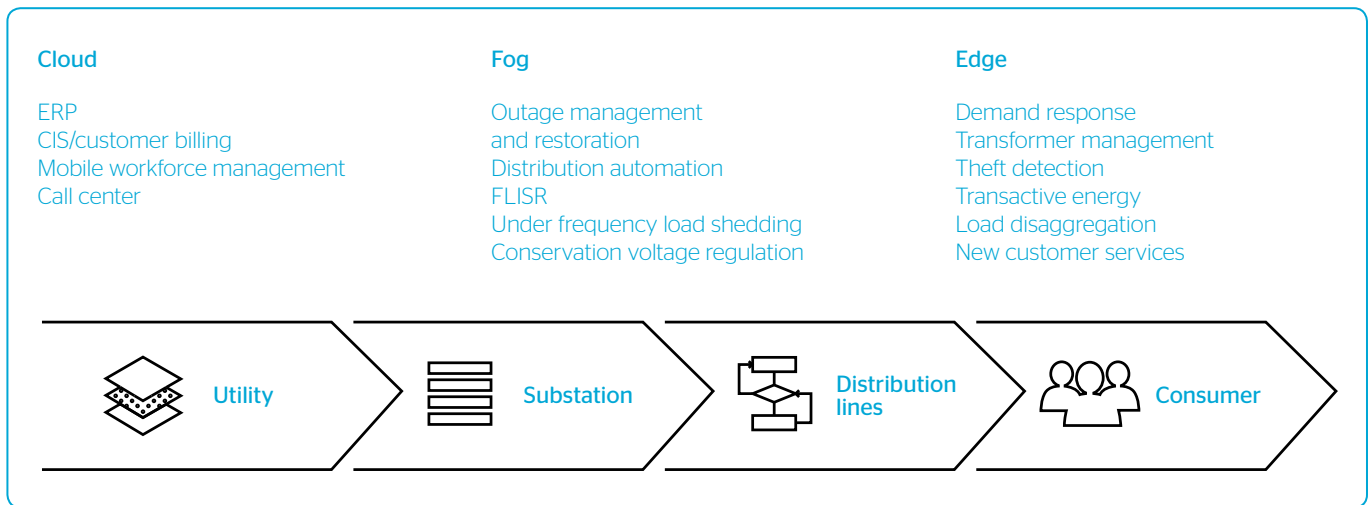
To cope with the vast amounts of data being generated by edge devices and sensors, utilities need to rethink their data management strategy. Certain types and amounts of data can now be processed and analyzed at the edge or at the network level rather than transported to a centralized or cloud data center. But this new approach to grid operations and data management also presents new kinds of risks. Adding connected intelligent devices in remote locations without regular inspection increases cybersecurity risks, as many of

these devices are not yet equipped with traditional IT protocols.

Beyond the business objectives, cross-functional teams must develop a new generation of grid management use case definitions and requirements to identify the sequence of steps – from data collection to analysis to action – for a given set of circumstances. It requires the development of new algorithms that balance localized, automated actions with centralized visibility and control.

Like the larger IoT movement itself, a technology revolution at the grid's edge is underway and gathering speed. Consumers are fueling this revolution with accelerated adoption of solar energy, EVs, smart home technologies and smart appliances, and now energy storage. Utilities that fail to leverage their incumbent status as an enabler of this movement will imperil their future growth opportunities.

Utility Applications Associated with Cloud, Fog and Edge



Why Atos for Digital Utilities

Atos integrates IT and OT to deliver real-time industrial IoT solutions for energy and utility companies. We work across the power, water, oil and gas industries – from production and distribution to transportation and retail services. With more than 35 years of utilities experience, over 3,000 industry specialists and an innovation-focused R&D culture, we help energy and utility companies drive digital change to realize business value across their organizations.

Learn more at <https://atos.net/en-na/lp/iot-utility-2019> or send an email to info.na@atos.net.

About Atos

Atos is a global leader in digital transformation with 120,000 employees in 73 countries and annual revenue of over € 12 billion.

European number one in Cloud, Cybersecurity and High-Performance Computing, the Group provides end-to-end Orchestrated Hybrid Cloud, Big Data, Business Applications and Digital Workplace solutions through its Digital Transformation Factory, as well as transactional services through Worldline, the European leader in the payment industry. With its cutting-edge technologies and industry knowledge, Atos supports the digital transformation of its clients across all business sectors. The Group is the Worldwide Information Technology Partner for the Olympic & Paralympic Games and operates under the brands Atos, Atos Syntel, Unify and Worldline. Atos is listed on the CAC40 Paris stock index.

Find out more about us

atos.net

atos.net/career

Let's start a discussion together



For more information: info.na@atos.net

Atos, the Atos logo, Atos Syntel, Unify, and Worldline are registered trademarks of the Atos group. April 2019. © 2019 Atos. Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.