White paper
Expert Community

Edge computing

Thought Leadership

Atos

Executive Summary

This whitepaper provides an overview of edge computing technology, use cases and existing market developments. It elaborates on analysis of the existing market status building a classification which entails consumer, software, hardware and born at edge categories. In addition, it relates to on-going initiatives and opportunities for further growth being explored by Atos, including the introduction of the Atos Edge Server.

Tables of contents

Edge computing: definition, motivation and drivers

Motivation

The unstoppable development of Internet of Things is expected to exponentially grow the number of connected devices worldwide. According to the latest estimates, these could reach more than 20 billion by 2020[1]. In this context, the challenge is not only the management of these things on and across the various platforms, but also, how to cope with the volume of data and content that all these connected things will produce. Researchers estimate that in the last two years we have created 90% of all available data on the internet, and this is expected to grow at 40% pace in the coming years due to the explosion of the number of available connected devices[2,3].

Following the pattern we have today in the existing IoT set-ups, generated sensor data would be transmitted over the wide area network in order to be centralized, processed and analyzed - creating additional volume of (enriched) data. This data, in addition to analytical models, are supposed to trigger actions either on the thing itself or in the upstream business systems – or even in other platforms that can make use of the data out of their original context (e.g. a car insurer using in-vehicle data to support a pay-as-you drive offering).

The number of device connections, the volume of data, the latency across different locations and networks, together with the asynchronous nature of many connections between data flow and analytical cloud services, make an existing IoT approach unsustainable in the long term. Besides, today's heterogeneous networks are not yet ready for managing the expected incredible growth in the number of end-points and the volume of data to be transferred, for which smart ways to distribute the data are not yet available.

Apart from the previous considerations, production critical environments need yet to address requirements for immediate response times. Nowadays, in many environments, the stability of the connection to centralized systems, as well as the latency incurred in the communication, make it impossible to achieve the necessary response times required to trigger and execute the actions in closed loop environments.

As part of addressing this issue, there is the pressing need to move information close to the source, transferring analytical model execution near to the sources of data, so to provide compute capability inside an environment where connectivity and response times can be controlled.

To address the above points a new class of edge connectivity and computing has emerged. It deals with network and local device management, pre-processing of the message streams (e.g. data can be pre-processed, aggregated or filtered and distributed) to ensure that the subscribers of a specific data stream only receive the data that is relevant for them, and that even the execution of analytical models and resulting actions is produced swiftly and in a distributed model.

1. Gartner, Leading the IoT, https://www.gartner.com/imagesrv/books/iot/iotEbook_digital.pdf
2. Big Data, for better or worse: 90% of world's data generated over last two years, http://www.sciencedaily.com/releases/2013/05/130522085217.htm
3. IDC, The digital Universe of Opportunities: Rich Data and the Increasing http://www.emc.com/leadership/digital-universe/2014iview/executive-summary.htm

What is edge computing?

Edge computing locates computing and storage resources at the edge of the network, with the intention of "getting data and computation at the right place and right time in a more decentralized manner"[4] and to elude non-essential data transmissions over the network. Gartner defines edge computing as "solutions that facilitate data processing at or near the source of data generation". At time of writing there still remains in the market significant confusion with regards to what is Edge computing and specifically which are the features and capabilities it offers. Diverse approaches exist, coined under diverse terms: Edge computing[5] is widely used; but also Fog Computing[6], as named by Cisco; and Multi-Access (Mobile) Edge Computing[7], focusing in RAN (Radio Access Network) technological aspects. It is important to note that this whitepaper we take the approach to Edge Computing as defined by the Linux Foundation Open Edge Computing Glossary[8] which defines the Edge Cloud as "Cloud-like capabilities located at the infrastructure edge, including from the user perspective access to elastically-allocated compute, data storage and network resources. Often operated as a seamless extension of a centralized public or private cloud, constructed from micro data centers deployed at the infrastructure edge."

A common topology for an edge computing installation is composed of three-layers. From bottom to top these are:

- **IoT devices:** IoT devices are connected to an edge device. IoT devices communicate via diverse communication protocols with the edge environment acting as data sources.
- **Edge nodes:** they enable data processing close to sources of data through near real-time data analytics and model execution. It offers divers communication and messaging protocols for data acquisition from near-by IoT devices and acts as temporal data storage.
- **Cloud services:** they develop management functionalities for both edge and IoT devices and they perform long-term data storage and analytics. Additionally, they provide the point for integration with other enterprise systems. In some advanced scenarios, such us autonomous vehicles, the limits among these three layers blur, the car being a rich environment of IoT devices and edge capacity all integrated in a single device.

The recognized advantages of edge computing are:

- **Privacy:** To avoid sending all data to be stored and processed on cloud servers.
- **Connectivity:** To avoid and reduce costs associated with streaming/sending all raw data to cloud services.
- **Latency and reliability:** To improve responsiveness and reliability by maximizing processing at the edge and thus minimizing dependence on cloud connection.

Edge and cloud manage processing, networking and storage resources. In this regard, they share essential mechanisms and attributes (for instance virtualization). Nevertheless, edge computing specifically focuses on addressing latency issues detected in large IoT scenarios and this brings a number of important specific characteristics to edge computing[9]. Among them:

- **Wide physical distribution:** in contrast to centralized cloud, the services and applications in edge computing are decentralized and distributed by nature.
- **Need for near real-time interactions:** edge computing aims to enable IoT actuation scenarios by bringing computation close to data sources. In doing so, there is a strong requirement for immediate response times.
- **Autonomy:** edge computing installations are not subject to data center boundaries and well-known fault tolerance practices. Therefore, they are prone to circumstances that can affect their availability and connectivity.
- **Heterogeneity.** Heterogeneity is a key characteristic of Edge computing which is present both in the variety of devices which form and connect to Edge computing instances, but also in the diversity of network and security protocols they are exposed to.
- **Interoperability and federation.** Typically, an Edge installation works together with a federated Cloud offering. For this, both Edge and Cloud environments inter-operate for operations such as Edge device management or long-term data storage.

Figure 1: Edge computing layered stack

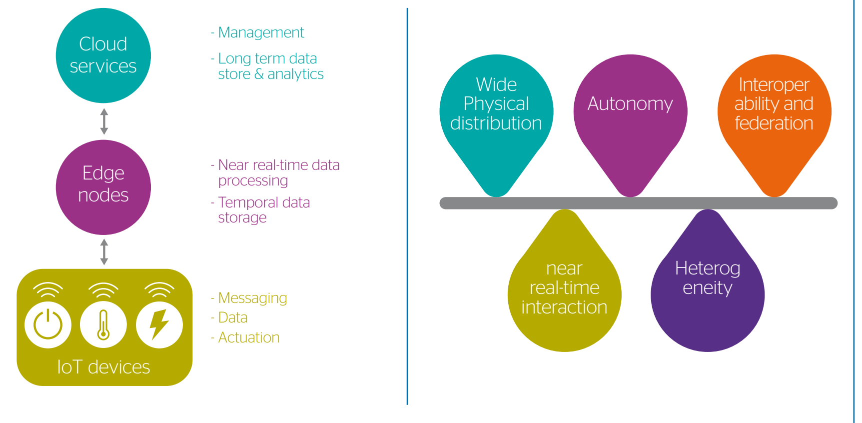

4. Bonomi, F., Milito, R., Natarajan, P., & Zhu, J. (2014). Fog computing: A platform for internet of things and analytics. In Big data and internet of things: A roadmap for smart environments (pp. 169-186). Springer, Cham.
5. FL Edge, https://www.lfedge.org/
6. https://www.openfogconsortium.org/
7. Multi-Access Edge Computing, https://www.etsi.org/technologies/multi-access-edge-computing
8. https://github.com/lf-edge/glossary
9. YM. Yannuzzi, R. Milito, R. Serral-Gracià, D. Montero and M. Nemirovsky, «Key ingredients in an IoT recipe: Fog Computing, Cloud computing, and more Fog Computing,» 2014 IEEE 19th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), Athens, 2014, pp. 325-329.doi: 10.1109/CAMAD.2014.7033259

Use cases

Edge use cases are mainly related to those where real-time (or almost real-time) reaction is needed and the latency to move to central cloud services does not ensure the necessary velocity on reactions. So, we may identify as key use cases the following:

- **Autonomous driving:** Tesla is one example of using computing in the edge but also centralized ML practices to improve the overall behavior of edge devices. We see examples in all transportation subsectors (cars, trucks, ships, drones, etc).
- **Process monitoring and operations analytics:** the monitoring of business process and operations through the extensive use of sensors and information systems are making inroads to increase efficiency of operations. For instance, we may think of refrigeration checks for food security and safety, where conditions are verified and warnings issued (or actions taken) if values fall below specific targets. Systems can be end user devices (personal fridges) or industrial ones. Actions may range from a request for additional stock when empty or for support assistance if values point to errors.
- **Asset tracking and monitoring (including predictive maintenance):** this has been a focus area so far, aligned with Industry 4.0 initiatives. The fact that machines equipped with sensors can be monitored in real-time. Together with the amount of data available to use with ML algorithms to learn from patterns and behaviors, allow for real-time monitoring to prevent failures and even, being able to forecast them therefore reducing the cost of down time, waiting for spare parts and repairs.
- **Employee, customer and citizen monitoring:** the use of smartphones, wearables and tagging devices with their related beacons (Bluetooth, NFC) and an increasing number of image recognition systems are putting the monitoring of persons (and reacting to that) center stage of new edge computing applications. Examples include using person/Face Recognition Systems, Human Behavior analysis (through image and location analysis), customer experience monitoring (and influencing though targeted advertising and notifications). Keeping in mind privacy concerns, there are also a wealth of uses around safety and compliance that will be a major push in this area.

These use cases are sometimes mixed to create business solutions that fill all requirements that customers may have in today's complex scenarios. As an example, we may think of healthcare and, in a more concrete manner, a hospital. We may have the asset tracking scenario to check the health and status of all equipment but also receive sensor information from patients and have some decisions being taken on the edge to improve customer care. All this increases the resilience of the overall structure having all actions and resources being used at the proper time and location depending on priority, latency and availability.

Edge Market Status

Edge computing enabling technologies are a major opportunity for existing IT stakeholders and new emergent ones. Gartner estimates that, "Currently, around 10% of enterprise-generated data is created and processed outside a traditional centralized data center or cloud. By 2022, Gartner predicts this figure will reach 50%." In terms of market size, analysts measure the market growth as, "6.72 USD billion by 2022 at a Compound Annual Growth Rate of 35.4%".

The edge market is still in its infancy, however its complexity is starting to be recognized. The edge computing market will allow diverse market actors to have their say in providing edge computing products and services. These range from system integrators, major cloud providers, hardware and software vendors, telecom operators and more.

In this scenario, different players and strategies come into play. With the objective to provide an overview of existing market developments we have structured our analysis of the edge market in the following categories:

- **Consumer edge:** represents vendors in the field of consumer electronics who are likely to address the opportunities that edge computing brings by embedding (or allowing to interact) edge intelligence in their devices.
- **Software edge:** these are vendors, mainly major cloud providers, who offer software platforms to be used in conjunction with their IoT cloud services and installed in edge devices provided by the user.
- **Born at edge:** while all the previous categories refer to existing vendors who are - in different ways - extending or adjusting their existing products or services to cover the edge market spectrum, here we refer to specialized edge computing vendors whose offerings only focus on this market.
- **Hardware edge:** in this category we present actual hardware device vendors who are including, as part of their products, the development or packaging of software tools that enable edge workload management. Atos is classified under this category.

For some of these categories we will be presenting examples of vendors to show how the edge opportunity is being addressed. Note that this analysis will present information available at time of writing (July 2018) and does not aim to be an exhaustive market guide instead we aim to provide exemplary cases for the established categorization.

Figure 1: Edge computing layered stack

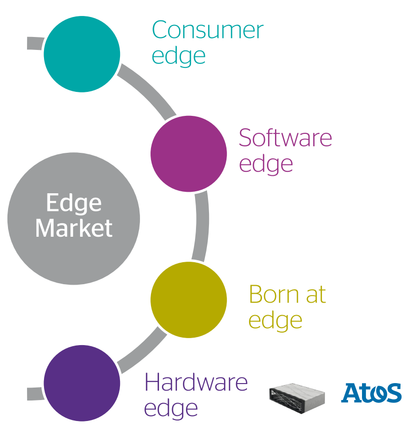

[10] Gartner Blog, "What edge Computing Means for Infrastructure and Operations Leaders", **https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/**
[11] Markets and Markets, "Edge Computing Market worth 6.72 Billion USD by 2022", **https://www.marketsandmarkets.com/PressReleases/edge-computing.asp**
[12] Gartner, "Edge Computing Challenges Go-to-Market Strategies in IoT", **https://www.gartner.com/doc/3877186/edge-computing-challenges-gotomarket-strategies**

The consumer edge

Edge can scale from personal devices (i.e Fitbit) to high performance compute in wind turbines and complex manufacturing (i.e Siemens). It can include B2B and B2C and even C2C scenarios. Edge has the potential to add value to almost all devices, but success is about more than internet connectivity and capturing data. The point is that the edge device allows capture and analysis of data with the goal of providing insight.

Even now embedding intelligence into devices is generally missing. Only a few market developments are starting to tackle these opportunities. General Electric is embedding edge computing into industrial internet and providing an edge to cloud distributed intelligence through its Predix platform[13]. Another recent example is provided by Amazon, that is reported to be building specific AI chips for its Alexa platform[14].

A very significant field of development for the future are autonomous vehicles for which Atos is developing crucial technology in the context of the European Processor Initiative[11].

The software edge

Major cloud vendors such as Amazon Web Services and Azure are taking the edge computing opportunity to extend their cloud services and to provide an additional entry point for their consumption.

Azure IoT edge[15] is available as Open Source platform[16]. It can execute both in Windows and Linux powered devices. IoT edge modules run as Docker containers. Its runtime provides monitoring and workload execution functionalities at the edge. It allows data pre-processing on-premises before sending it to Azure Cloud services. Azure services, which can run on edge devices include Azure Machine Learning, Stream Analytics Azure Functions, AI services and the Azure IoT Hub.

AWS Greengrass[17,18] software stack is available for both ARM and x86 devices with minimum required capacity. It also takes into account communication and management for micro-controllers (Amazon FreeRTOS). AWS Greengrass is a Serverless computing platforms-based AWS Lambda which also provides: data messaging and edge to cloud sync via interoperability with AWS IoT Cloud services. Similarly to Azure, it has recently further developed ML inference capabilities in Greengrass ML. These ML models, which have been developed and trained in the cloud, can be deployed and executed locally in the edge device.

In either case, the major cloud players bring the opportunity to have a unified ecosystem, but potentially locked in to that vendor. Running the software as containers on top of Linux allows easy deployment, but the code written may well end up locking you in to that vendor longer term without significant effort. Whether using Azure's IoT edge and IoT Hub or AWS Greengrass, you can now allow edge distributed logic to be written in the same format as what can run centrally, easily allowing a centralized model to move to an edge deployment at a later stage. Serverless code, whether it be Azure Functions or AWS Lambda, allows local decisions to be taken rapidly and even insight and analytics of data, with low latency, even in scenarios where internet access might be down or intermittent, while at the same time allowing aggregation of data back to the central cloud for global scale reporting.

The ability to train a ML model rapidly in the cloud and then roll that model out locally to thousands of distributed locations empowers the edge. Models that even a few years ago would seem impossible to run on the edge suddenly become feasible on small and low power hardware, both in generic processors, but also potentially in ASIC based hardware.

Taking it a step further, Azure Stack brings another potential to run considerably higher power applications, bringing the power of Azure Cloud in to your own data center with hundreds of cores and terabytes of memory for local delivery of Azure native services, whether it be for latency, geographic availability or data sensitivity purposes. The key is the same models from public Azure have the potential not only to be moved to smaller lower power hardware, but also to distributed scenarios offering huge levels of local processing power.

13. GE, Predix Platform, https://www.ge.com/digital/predix-platform-foundation-digital-industrial-applications
14. The Verge, "Amazon is reportedly following Apple and Google by designing custom AI chips for Alexa", https://www.theverge.com/2018/2/12/17004734/amazon-custom-alexa-echo-ai-chips-smart-speaker
15. European Commission, "European Processor Initiative: consortium to develop Europe's microprocessors for future supercomputers", https://ec.europa.eu/digital-single-market/en/news/european-processor-initiative-consortium-develop-europes-microprocessors-future-supercomputers
16. Azure IoT edge, https://azure.microsoft.com/en-us/services/iot-edge/
17. Azure IoT GitHub. https://github.com/Azure/iot-edge/
18. Amazon Web Services Greengrass, https://aws.amazon.com/greengrass/

The hardware edge

While vendors in this category have traditionally focused on devices and equipment, increasingly they are trying to embrace both hardware and software stacks, so as to build overall edge solutions as part of their value proposition. Examples of this approach are Cisco and Dell.

Cisco's was a precursor in identifying the value for computing at the edge, coined under the Fog Computing term[19]. Cisco sees the need for a handling the volume, variety, and velocity of IoT data in a new computing model. Recently,

Cisco has launched Cisco Industrial Compute Gateway IC3000[20]. This is equipped with its Kinetic edge and Fog Processing Module enabling a software stack to distribute computing. In addition, the Cisco IOx application framework offers security capabilities as well as consistent management and hosting across network infrastructure products, including Cisco routers, switches, and compute modules.

Dell edge Gateway Series[21] is purposefully built for developing industrial automation. These Gateway Series offer specific IoT hardware equipped with Intel base Atom processor enabling integration with diverse protocols (ModBus, BACnet and ZigBee) as well as, network access means (Wi-Fi, WWAN and Ethernet).

Increasingly commodity cheap hardware such as Raspberry Pis and off the shelf x86 servers are a clear platform of choice to provide edge Gateway features, playing a significant role due to affordability and huge scale, which is achievable depending on the requirements.

The Raspberry Pi could be said to have started (or perhaps re-started) a computer revolution, bringing a new generation of users to 'home-brew' computer. The low cost, low power and commodity nature of the Pi has made a ideal starting point for producing commodity IoT edge devices. Both Azure and AWS include guides via Github on how to use these commodity devices and integrate them in to your IoT environment. Obviously, such a commodity device is not for every circumstance but has helped encourage a rapid growth of experimentation at the edge. Beyond this initial experimentation with Edge computing, progressively we observe in IoT applications not only the need of having rapid data processing close to the sources, but the crucial requirement to provide secure, intelligent monitoring and even, handling of the physical environment so to enable next generation, transformative AI and IoT applications.

Therefore, next generation Edge computing environments need to cope with the emerging necessity of offering AI inference and sophisticated data analytics at the Edge, which encompass steam video, image and audio near real time processing.

Born at edge

The development of the edge computing market is also bringing the emergence of specialized providers that offer targeted products and solutions.

For instance, FogHorn products specifically tackle the Industrial IoT market and centers its offering in its Complex Event Processing technology which is described as the "world's most compact, advanced and feature rich"[22]. This technology is provided in two different platform packagings, in standard and micro editions. Both editions include remote edge flee devices management including device and platform monitoring, configurability, as well as, automated deployment of custom, platform and analytics software. Differences among these are in their deployment memory requirements, less than 256MB devices for micro edition and 2G for standard edition. More in detail, Micro edition provides a C++ SDK for application development, data pre-processing, analytics, machine learning and visualization capabilities. Features available in standard edition include micro-edition's features with extended support for Industrial protocols, data access protocols, cloud interoperability and local data persistence. Cloud technologies or providers offering compatibility with this solution are not specified. FogHorn presents as a main novelty of its products edgeML technology, which aims to enable execution of machine learning models in edge constrained devices.

Another example is the French company Tell Me Plus. This company has been chosen as part of Microsoft's "Accelerator Program", where Microsoft looks to equip participants with tools to help them develop their product and grow their organization. Similarly to FogHorn, Tell Me Plus software platform, predictive Objects, targets the Industrial IoT sector.

Other novel general-purpose edge platform vendors are Crosser and Rigado.

Besides, we also start to encounter niche offerings which are addressing specific needs of edge solutions. Examples of these are: IoTium which focus on edge security aspects; and Wirepas that addresses the needs of IoT networking at scale.

19. Amazon Web Services Greengrass FAQs, https://aws.amazon.com/greengrass/faqs/
20. CISCO, "Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are", https://www.cisco.com/c/dam/en_us/solutions/trends/iot/docs/computing-overview.pdf
21. CISCO, "Helping Enterprises Scale IoT Deployments with Secure Computing at the edge", https://blogs.cisco.com/digital/helping-enterprises-scale-iot-deployments-with-secure-computing-at-the-edge
22. Dell edge Gateways for IoT, https://www.dell.com/en-us/work/shop/gateways-embedded-computing/sf/edge-gateway
23. FogHorn products, https://www.foghorn.io/products

Challenges for adoption of edge

While the edge market is rapidly developing and benefits of edge computing relate to improved applications performance derived from reduced latency in data processing, it comes with a series of challenges related to its adoption:

The first relates to the novelty of the concept and lack of awareness of the benefits. While the edge technology is still under development and not fully consolidated in the market, users have difficulty in getting concrete figures and information on performance gains and applicability to its concrete usage scenarios.

In relation to this, there still remains confusion in some instances about the complementary or substitutivity between edge and cloud. Whereas common perception is that edge computing completes cloud computing, offering an alternative solution for which existing services in the cloud are not sufficient, the interrelations and interoperability among both approaches still appear uncertain to potential consumers of this technology.

Despite the fact that some initiatives such as OpenFog, ETSI and edge Computing consortium start to emerge in order to offer both certification and standardization schemes for this new technological environment there is not yet a clear understanding of required standardization efforts and vendor adoption in the market.

It is worth mentioning that edge technology also brings a number of important challenges not only to users to adopt technology but to the commercial providers of these offerings such as service management in geographically distributed environments outside the data center limits.

Atos developments and opportunities

Atos, via it's cognitive analytics offerings, is already a significant player in the IoT marketplace. The Atos cognitive analytics tagline is «Turning Data into Business Outcome», which implies that we are supporting the full loop from the Edge device to collecting data from the device to trigger actions either in the device itself or in the associated business systems.

Edge computing plays an important role in any of these setups : it provides the means to connect the devices to manage the (in some cases bi-directional) data flow and associated actions, as well as, the volumes of data that are produced.

More strategically, Atos will launch its first Edge server, BullSequana Edge.

Atos is one of the few companies that can provide edge computing solutions at this level. The prototype description highlights the ability to provide AI computing power at the edge with rich security features. Moving up the stack - using edge devices that provide connectivity and device for their compute capability opens the door to move analytical execution close to the devices. Analytical models are being developed (often with the help of Machine Learning and Artificial Intelligence) on the centralized platform and the (sub-) models are distributed to be executed on the edge device so to either reduce the volume of data transmitted or to trigger automated actions. In complex business scenarios – e.g. remote factories relying on Manufacturing Enterprise Systems (MES) data – edge devices will start caching external data from the MES system, use real-time data from the shop floor which itself is analyzed on the edge device to provide fast response. The opportunity to build and integrate towards such scenario requires Atos capabilities across the board – IT/OT connectivity, Integration of shop-floor and top-floor data, analytics on both the central systems as well as on the edge and creation of micro-service or apps on the edge itself to provide decision support or to trigger actions.

From a connectivity and data management side, Atos own edge solutions («cognitive analytics Connectivity Platform» (CCP), «Cloud Industrial Supervision» (CIS)) help us build platforms to connect devices (like sensors on the shop floor, vehicles, power valves) and help distributing compute and analytical tasks. These are software solutions built on top of hardware, assuming end-points are available. CIS is an Atos cognitive analytics IoT component that actually helps to control the flow in such distributed scenarios.

With regards to long term development of edge computing, Atos is working with its Bull subsidiary on edge devices that provide massive compute and analytical capabilities (incl. Machine Learning and Artificial Intelligence) - this is being announced as «European Processor Initiative with Deep Learning acceleration» (EPI) for future availability.

At the level of software stack, Atos Research and Innovation is participating in a set of related H2020 initiatives taking key aspects of future edge computing developments: AGILE project[24], which builds an IoT gateway that offers data and workload portability with key public cloud offerings; DITAS project[25], that develops and implements the concept of Virtual Data Containers and specifies data movement approaches across edge and cloud environments; and mF2C project[26], which focuses on providing across edge and cloud workload orchestration techniques.

[24] Agile Project, **http://agile-iot.eu/**
[25] DITAS project, **https://www.ditas-project.eu/**
[26] mF2C Project, **http://mF2c-project.eu**

Conclusion and call to action

Atos is in a privileged position in order to build a compelling edge computing offering which takes advantage of the uniqueness and diversity of Atos capabilities.

Closely related to the definition of its IoT strategy, edge computing offering can be a key development for the coming years for Atos by agglutinating and exploiting a series of capabilities that makes it exceptional in the market: ability to deliver its own edge compute servers cutting-edge IoT security assets; capacity to build a software platform with robust analytics and machine learning inference features and interoperability with own and major cloud providers services (as Atos cognitive analytics evolution); as well as, detailed sector specific knowledge and consultancy skills to bring this novel offering to the market.

It can be observed from current market developments detailed in Section 2, that major cloud providers are focusing on provision of software platforms, relying on a network of hardware providers on which to make their offerings compatible. In the long term, these providers won't be able to adapt these to the specific needs of their software platforms. In addition, it is important to note, these providers are offering their customers a not negligible vendor lock-in (even at programming model).

Hardware providers which are strong and capable in developing and deploying devices at scale, are trying to build software platforms from scratch, and are not yet embracing possibilities of building partnerships with cloud vendors for this purpose. Therefore, to some extent, they are lagging behind at the level of the software capabilities they offer to customers having to build applications. Although the use of commodity cheap hardware is increasingly becoming popular, it has not yet had a reference Open Source edge computing project that delivers a software stack which helps to build a complete solution beyond pilot execution.

At the same time, the novel market actors are focusing on the industrial internet arena which nowadays concentrates the majority of edge commercial efforts. The list of niche products available today is significantly limited talking solely security and networking.

From this analysis it becomes clear that there is not yet a provider able to offer an edge computing end to end solution comparable to Atos' potential in the field. Atos is able to build on top of horizontal assets (hardware platforms, software security, analytics and interoperable cloud software platforms) specific customizations to vertical sectors and IoT platforms which could, complemented with domain specific knowledge and consultancy services, facilitate adoption paths for customers of this future crucial technology.

Authors

Ana Juan Ferrer	Head of Next Generation Cloud Lab at Atos Research and Innovation, Atos Distinguished Expert in Cloud Domain and member of the Atos Scientific Community
Adam Dolman	Head of Public Cloud for Atos Global IDM ESO, a member of the Atos Scientific Community, and an Atos Expert in the Cloud Domain
Akshay Bagewadi	Cloud Expert & System Consultant Lead at Atos
Alex Caballero	CTO IDM Iberia, Atos Senior Expert in Cloud Domain, and member of Atos Scientific Community
Claude Deure	HPC & Big Data Marketing, BDS Global
Darren Ratcliffe	Distinguished Expert & Cloud Domain lead, Head of Technical Strategy within Business & Platform solutions
Gary Burt	Atos Distinguished Expert, Cloud Product Manager and Hybrid Cloud Business Strategist
Gauvain Girault	CTO Big Data, BDS Global
Gert Prieber	BullSequana Edge Product Manager, BDS Global
Klaus Ottradovetz	Atos Distinguished Expert Cloud & IoT platforms, CTO office on cloud and IoT platforms and technologies
Prashant Chamarty	Expert Sales at IoT and Cloud Centre of Excellence, Atos Senior Expert in Cloud Domain
Ripal Kaila	Sr. Expert Application domain, Sr. Solution Architect Cloud, SAP HANA
Roy White	Global MS Digital Cloud, Infrastructure & Applications
Sameep Bhate	Solution Architect

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: tessa.david@atos.net

CT_190425_MA_ExpertCommunity-EdgeComputing_WP