Edge Intelligence: The Central Cloud is Dead – Long Live the Edge Cloud!

Authors:

Yun Chao Hu, Huawei Technologies Duesseldorf GmbH**,** Yunchao.Hu@huawei.com

Andreea Ancuta Corici, Fraunhofer Institute FOKUS, andreea.ancuta.corici@fokus.fraunhofer.de

Mike Cronin, Johnson Controls Innovation Garage, mcronin@tycoint.com

Marc Emmelmann, Fraunhofer Institute FOKUS, marc.emmelmann@fokus.fraunhofer.de

Bing Liu, Huawei Technologies Co, Ltd, remy.liubing@huawei.com

Kunlun Lu, Huawei Technologies Co. Ltd., lukunlun@huawei.com

Thomas Magedanz, Fraunhofer Institute FOKUS, thomas.magedanz@fokus.fraunhofer.de

Tetsushi Matsuda, Mitsubishi Electric Corp., matsuda.tetsushi@dh.mitsubishielectric.co.jp

Andreas Vogel, SAP, a.vogel@sap.se

Victor Kueh, Huawei Technologies UK CO., LTD, Victor.Kueh@huawei.com

Peter Lanctot, IEC, Regional Center for North America, pjl@iec.ch

INTRODUCTION

Edge intelligence allows bringing data (pre-) processing and decision-making closer to the data source, which reduces delays in communication. In addition, such (pre-) processing makes it possible to accumulate and condense data before forwarding it to Internet of Things (IoT) core services in the cloud or storing it, which perfectly matches the capacities offered by the upcoming fifth generation wireless technology (5G) networks providing localized throughput and delay enhancements. Edge computing makes processing and storage resources available in close proximity to edge devices or sensors, complementing centralized cloud nodes and thus allowing for analytics and information generation close to the origin and consumption of data. Supplementary resources may even reside on end devices that might not be continuously connected to the backbone network. Additionally, edge intelligence allows future applications to depend on context awareness capabilities for mutual detection and proximity services, (near) real-time responsiveness for a tactile Internet, data analytics at the edge and/or end device and device-to-device communication capabilities.

As processors, microcontrollers and connectivity modules are embedded into a plethora of new devices, the application of edge intelligence in smart appliances, wearables, industrial machines, automotive driver assistance systems, smart buildings and the like continues to increase. To enable and realize IoT's true value, the trend toward adopting edge intelligence, which pushes processing for data-intensive applications away from the cloud to the edge of the network, continues to expand.

TREND DRIVERS AND STATE OF THE ART FOR EDGE INTELLIGENCE

This article analyses manufacturing, smart building, asset management, smart grid and transportation industrial verticals. Some of the most stringent needs of the analyzed industries that can be solved through placing intelligence on the edge computing units include:

- Mobility: industries demand in terms of networking support (mobility and wireless broadband) is a high degree of mobility, especially in terms of handovers. At the same time, the "quality of service" and session handover management are critical aspects that can benefit from intelligence in the network components.
- Ultra-low latency in decision-making: decisions on detection or actuation have to be taken within a delay of less than tens of milliseconds. For this, intelligence residing at the edge can help lower the delay and achieve the targeted response time.
- Autonomy: a key requirement for use cases is autonomous to continuing operation without connection to core server or service, to prevent damage to persons, goods or infrastructure.
- Security: security is a feature that can never receive too much attention. Access control to physical or virtual resources (e.g. data) has to be ensured. Locally provisioned or learned policies and other mechanisms for

Need	Industry sector			
	Manufacturing	Buildings & Facilities	Energy & Utilities	Consumer & Home
Mobility	55	10	10	55
Ultra Low latency (<10ms)	95	85	5	15
Autonomy	95	100	100	50
Security	100	100	100	25
Local Network Bandwith	100	90	10	35
WAN Network Bandwith	35	55	10	55
Peer-to-Peer Communication	80	85	50	90
Prioritization	100	15	90	10
Self-organization, discovery	60	20	40	65
Artificial Intelligence/ Machine Learning	100	100	85	45

Table 1: High level needs for edge intelligence by industry sector

authentication and authorization running on edge computing units are envisioned to enable fast adaptability of the systems.

- Prioritization: refers to prioritizing the communication information depending on the type of included data (data gathering, alarms).
- Self-organization, discovery: discovery of the capabilities of the devices and services and their role in the infrastructure, so that operations can be handed over from humans to intelligent software.

Artificial intelligence/machine learning: may require and can gain from algorithms that automate decisions or alarms without human intervention. Adaptive software that analyzes data on edge computing units can deliver actions and awareness with minimal latency. A broad view of industry needs on different capabilities of the edge intelligence is summarized in Table 1 using a heat map to suggest importance for a specific industry. The scale ranges from 0 to 100, with 0 being unnecessary and 100 value in case the feature is vitally important for the industry. The assignment of the values is done based on a discussion between the authors of the article, having as background the previous projects.

Hardware Evolution

The trend within data centers is to shift to smaller, more agile (i.e. movable) data centers deployed towards the edge. This includes, for example, “edge caching”

approaches, adopted by companies such as Google to minimize latency in response. Others use scaled-down "out of the box" data centers, co-located or situated near to ISP nodes. This can be taken to its logical conclusion, if such caching/hosting is implemented at the base station level in wireless networks. Finally this trend can be extended to allow containerization at the base station and thus enable the docking of third party applications there.

As billions of endpoint devices need to connect, a critical component of future Internet of Things systems "IoT gateway." New generation IoT gateways are opening up huge opportunities to push processing closer to the edge, improving responsiveness and supporting new operating models. The IoT gateway serves as an important bridge between operations and IT and also provides a cost-effective business model by making use of lightweight storage, networking, and computing that can operate in industrial environments. By adding IoT gateways, the current field deployment could require no change in order to run new applications such as predictive maintenance on the gateway.

In scenarios such as smart manufacturing, with more and more robots, CNC machine tools, and the like, generating massive real-time data in the field, greater computing and storage resources will be necessary. In such situations a local cloud at the edge represents a good choice. More edge servers will be interconnected and provide pooled and scalable resources. Several software programs or services can be deployed on an edge node simultaneously.

Software Evolution

Most of the IoT Edge Computing Node technology, runs on flavors of the Linux or Microsoft Windows (using the .net platform) operating systems using various processor architectures. An industry-wide trend is emerging to package edge computing capabilities into micro services and deploys them within containers on IoT edge computing nodes, as illustrated in Figure 1. Containers are user space instances executed in an operating system kernel providing strong isolation between them. The operating system kernel can provide resources management between different containers. Containers provide security though isolation; they also serve as deployment units that simplify lifecycle

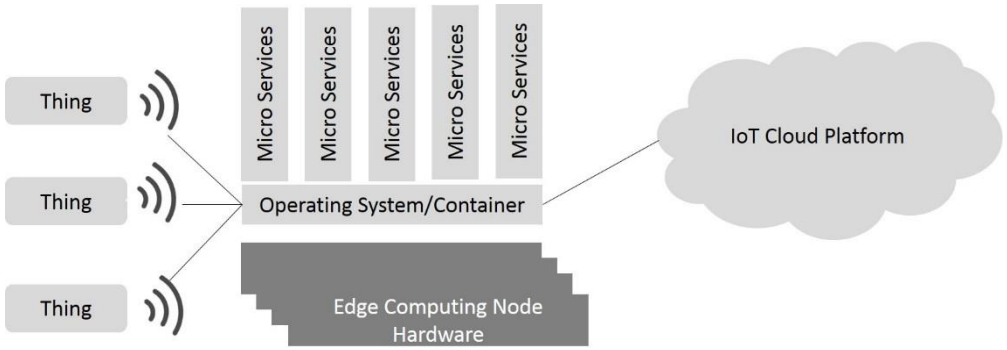

Figure 1: Edge computing

management through less interdependency and complexity.

An exception to this Linux predominance is found in the manufacturing industry, where various versions of the Windows operating system and Microsoft's .NET framework make up the majority of implementations.

Containerization and Micro-services

A key architectural goal for edge computing should be to isolate services that can be deployed anywhere. In this case, "anywhere" can be on a smart device, on an IoT gateway, in a micro data center, inside a telecommunications network, in a data center, or in a public or private cloud. This isolation can be achieved through self-contained micro services or pods (a group of services sharing resources, e.g. a database).

With the expansion of end and edge devices, deploying and maintaining code bases will become more difficult. Containers can help mitigate this additional overhead by providing virtualization on the operating system level..

Unlike virtual machines, containers impose little or no overhead. However, containers are less flexible than virtualization, only one operating system is used, Linux being the most popular one.

Docker has been the most popular container technology, but more recently "rkt" and Rocket have become widely used. The Open Container Initiative (OCI), a Linux Foundation project, aims to bring the competing container technologies together.

Kubernetes also introduces the concept of pods, which is a group of services deployed in containers that share certain resources, e.g. a data base, and interact via inter-process communication. Pods are an interesting concept for edge computing as they allow services to share scarce resources. Kubernetes, another open source project, is an open-source project by the Cloud Native Computing Foundation (CNCF) for automating deployment, scaling and management of containerized applications.

While containers are the preferred technology for deploying pods and micro-services in the cloud, they are also very applicable to edge computing and early adopters are already using containers or container-like technology on the edge. This allows the deployment of the same micro-services in the cloud, on premise and on the edge without touching the code.

It would be beneficial to edge computing to provide a uniform runtime infrastructure to deploy services both in the cloud and the edge. By deploying containers at the edge, one can optimize applications and services for bandwidth, storage, and compute power to meet the latency, scale, reliability, and ultimately cost requirements.

A commonly accepted open source platform would reduce total cost of ownership for customers and would allow solution providers to focus on

providing value-adding services and industry-specific solutions, not on navigating different platforms.

Machine Learning

Many applications of machine learning exist today, and the number is growing rapidly. There are four types of machine learning:

- supervised learning (also called inductive learning): training data includes 'all' desired outputs;
- unsupervised learning: training data does not include desired outputs; an example is clustering;
- semi-supervised learning: training data includes a few desired outputs;
- Reinforcement learning: rewards result from a sequence of actions; this appeals to AI practitioners, it is the most ambitious type of learning.

Supervised learning is the most mature, studied and by machine learning algorithms. Its popularity stems from the fact that learning with supervision is much easier than learning without supervision.

The next step beyond machine learning involves a complementary area called artificial intelligence (AI), which leans more on methods such as neural networks and natural language processing that seek to mimic the operation of the human brain.

The advantages of any supervised learning for IoT applications are clear – for sensor applications (including audio and video), the training data can be generated and models refined offline. Once refined, the models can be loaded at the sensor level. There are clear advantages to this approach:

- Core software updates are reduced – rules and operating code are separated and updates are limited to models only;
- Security is enhanced, an AI model or neural network is replacing a traditional rules based approach. So when this code is machine generated from a model (or decisions are implemented directly from the model) it is much harder to hack. The hacking problem becomes how to bias the data input to the model to try and "nudge" it's results to duplicate the incorrect behavior you wish to create instead of changing or intercepting the code. And when the model changes then the hacker has to start from scratch again.
- Model comparison can be performed at the CPU level, if the hardware supports it.

Network Evolution

In the domain of Core Networks, by adopting the paradigm to deploy network functions (NF) as programs on top of common off-the-shelf hardware, the main focus of the technology completely shifted towards the dynamicity offered by software programs, resulting in software networks. Being the most customizable form of control, software programs represent the full convergence between the telecom and IT industries unleashing new forms of innovation.

In the world of software-defined networking (SDN), the intelligence resides in the SDN controller that informs one or multiple SDN switches how to route packets (Figure 3-27). The SDN switches are programmable switches that interface via a protocol with the SDN controller. Monitoring of the data flows is also supported so that the SDN controller can be informed about the traffic the SDN switches are handling.

The flexibility of the software-defined networks is that initial switching algorithms or policies can be provided to the switching layer and, based on rules/policies/algorithms provided by the edge nodes, the master switching entity can adapt and improve the edge switching intelligence. The SDN switches can receive algorithms to process the data, and based on the traffic that is monitored can improve the algorithm and upload it to the SDN controller. Thus, SDN provides highly efficient and low-cost automatic Operation and Maintenance (O&M) and realizes policy collaboration, convergence of networks and security.

Network Slicing is a technology that allows Service Providers to build their own edge or core network using core network software components.

The 5G Core Network standardized by 3GPP takes advantage of the SDN as well as network slicing technologies, allowing for isolated dedicated networks to be built and managed. The networks can then adapt based on triggers coming from the application domains, e.g. alarms from detected security threats or other events.

Due to the softwarization of the 5G Core Network, there is a uniform approach for connecting end devices across the range of access networks from Wi-Fi and LTE towards NB-IoT, unlicensed spectrum LTE (LTE-U) and ubiquitous satellite connections.

USE CASES AND REQUIREMENTS FOR EDGE INTELLIGENCE

Factory Productivity Improvement

For factory operators, it is important to improve productivity to maximize profits and minimize waste. Big data analysis realized through a combination of edge computing and cloud computing can contribute substantially to factory productivity improvement. Edge computing with intelligence converts data collected from the shop floor into data which can be analyzed by cloud computing. The analysis result leads to reduction in number of products identified to be defective by mistake and also enables predictive maintenance, which can reduce factory down time.

Connected City Lighting

While bringing safety, convenience, and ambience to urban citizens, city lights consume a vast amount of energy and operational costs for municipalities. To address these issues, connected city lighting is being developed to optimize lighting coverage and resource usage. Firstly, integrated smart lighting policies ensure efficiency by controlling lights according to time, ambient light conditions, and weather. Secondly, online monitoring of each light's status in real time can automatically alert the maintenance through edge computing node once a malfunction occurs. The

traditional manual inspection is no longer necessary; therefore operational expenditures can be greatly reduced. Thirdly, if connectivity to the lights is lost, the systems can continue operating by using the policy from the edge device, and not rely upon the cloud or datacenter.

In the future, light poles will move beyond single functions (the lighting). Many other functional modules can be added, such as environment/utility monitoring, video surveillance, vehicle-to-infrastructure (V2I) communication devices, and so on, making the light poles become an integrated system of sensing and service provision.

Smart Elevators

Taller buildings and access make elevators indispensable in cities. The operation and maintenance of elevators is considerably expensive due to manual inspection, fault detection and repairs. Smart elevator with edge and cloud intelligence allow vendors to upgrade from inefficient, expensive preventive maintenance models to next-generation, real-time, targeted, predictive maintenance, extending value from products to services.

Hundreds of sensors are deployed to monitor the elevator's status. Based on this data, the edge computing node is capable of detecting potential device faults early and sending out the alarms immediately. When the edge computing node fails to connect to the cloud, the data can be stored locally until the connection recovers. By analyzing the historical data at the cloud, faults can even be predicted, so that maintenance is given accordingly before a fault actually occurs. New features of faults can also be extracted by artificial intelligence at the cloud and then downloaded by the edge computing node.

Indoor Location Tracking

The bandwidth requirements for indoor location tracking are moderate: approximately 2 MB, with very low latency (<1 ms) and low contention. The system requires a backhaul of trilateration data for a number of sensor sources (all normalized to IP/UDP packets) and conversion into a high quality location estimate. For high value asset tracking, real-time location computation require mathematical results and do not afford the delay introduced by communication to the cloud. As a result, having a gateway which processes the sensor samples as close as possible to the source, while maintaining the connection with the cloud, or at least outside the customer premises, are critical for a system of this type. Within the gateway module is an embedded, tunable, machine intelligence module to perform the location estimation, which then forwards real world positions and user status to the administrative/UI module in the cloud. The model may also be tuned and the Machine Intelligence (MI) module updated.

Lone Worker Safety

For lone worker safety an intelligent gateway is used to receive signals indicating the location of a particular employee. In such cases an MI module would also be embedded in the GPS signal transmitter, which would use an MI module to characterize the wearer's gait and orientation. The module would "learn" over a period the "normal" behavior of an individual, and thus be able to generate an

alarm should that behavior change due to accident, attack etc. Since decision-making is local, the transmission bandwidth would be very low and latency could be in seconds. The intelligent gateway in this case would function only as a data consolidator and could also have UI, etc., installed.

Physical Access Control – Tailgating Detection at Doors

Tailgating at security doors can be achieved by combining MI module and an inexpensive stereo camera, to ensure that only the right person got through secured doors. Badge readers(BLE/RFID for example) can add additional security,: i.e., the door opens as the person approaches, but only if a valid signal is read from each person (if there is more than one) approaching. If more than one person approaches, and fewer badge signals are read, then an alarm will be registered if they all enter the doorway. By using an inexpensive camera with a processing module, latency and bandwidth are reduced over sending to the cloud.

Fire Detection via Surveillance Cameras

An MI module is used with access to video streams incoming from non-specialized video cameras and scans those streams for fire traces using a trained model. In the case of a fire, latency is an issue, and the local processing can greatly increase response to minimize loss. When there are no incidents, bandwidth is minimized and only summary data is sent.

TECHNOLOGY GAPS

This section describes different technological gaps related to edge intelligence that need to be addressed to support requirements for the use described in the previous section.

Factory Productivity Improvement

(1) Credibility of information generated by edge computing

In order to realize edge intelligence, it is necessary to establish credible data, as the data are the raw materials for future analysis and decision-making. Now, there are technologies which can be used to ensure the credibility of data exist, such as technologies to secure the identity and the integrity of device-generating data and technologies to protect information being communicated over a network. In addition to enhancing technologies, systematic ways to guarantee credibility of the data by making use of such technologies and operational measures including physical security.

(2) Assisted/automatic optimization of system operation

In order to reduce the operational cost, it is necessary to optimize system operation and to adjust operational parameters more efficiently. At present, many productivity decisions in a factory are done manually by experienced engineers, and even the data for such decision-making can be collected. It would be desirable to make optimization easier by automating data collection and assisting engineers in the optimization analysis, ultimately automating optimization processes.

(3) Open platform for edge computing

At present, there are multiple vendors of edge computing products, and the edge computing (program execution) environment of those products is proprietary. This situation makes it difficult for developers of edge computing applications to develop a portable application which can run across multiple edge computing products. It would be desirable to utilize the wisdom of people from various disciplines in system development by enlarging the developer community, as in the case with Smartphone applications. Making the edge computing environment open can contribute to that purpose. Additionally, measurement such as certification of an application to guarantee that the application does not behave maliciously and is not harmful to a system, and lifecycle management of an application would need to be in place.

Smart City Lighting

(1) Function add and removal for Pole networking

When deploying lighting poles, the installed function modules should first be identified, and then neighboring devices and Edge Compute Nodes (ECNs) specified with which to establish connections. In this way, the network can be built without much human intervention. The networking topology respond to conditions such as connection qualities and the computation load/capabilities of ECNs. Once a function module is added or removed from a pole, the network and the applications at the ECN should change accordingly.

ECNs are machines deployed close to the end device of an infrastructure, having the role to take local decisions based on sensed or received information and policies or algorithms from the core servers. The decisions can be applied (physically) locally or communicated to the core servers for aggregation and final decision.

(2) Evolutionary lighting policies and energy management policies

The basic lighting policy is drawn up based on the time, date and geographical location. Some reactive policies can be made according to the environmental brightness and the proximity of vehicles and pedestrians. The edge and the cloud should have learning abilities to build continuously evolving rules. Since lamps and the other modules, especially the charging module, are part of the smart grid, the energy management policies should also be optimized based upon predicted use.

(3) Cloud offloading and privacy: process the data locally at the edge

The allocation of processes between the edge and the cloud must be defined: the time-critical data must be processed at the edge to get a timely response; some lower stage processing can be conducted at the edge, such as filtering and aggregation, so that the cloud can be offloaded; for privacy reasons, some data must be processed locally, e.g. in video surveillance, only abnormal events are reported to the cloud while the citizens' portrait should be protected.

(4) Interaction between ECNs, user information synchronization, control and orchestration

In some scenarios, the ECNs need to work in a cooperative way. For instance, at night, lamps need to increase their brightness when vehicles and pedestrians approach to improve safety and decrease the brightness when they leave to save energy. This process is expected to be continuous, especially the handover between two ECNs. Based on the headings, velocities and positions, an ECN can predict the next ECN that the vehicle will pass and remind the latter to get ready for handover.

Smart Elevator

In addition to cloud offloading and privacy, and interaction between ECNs functionality, as described in the smart city lighting use case, the smart elevator use case would require:

(1) Diagnosis and predictive maintenance

Elevator safety is paramount, and fault detection and maintenance is critically important. Currently, on-site maintenance requires highly trained personnel using advanced tool sets to diagnose and repair, which leads to extra expense in staff training and equipping. Moreover, it often takes a considerable amount of time to identify the fault and find the corresponding solution.

Autonomous, accurate and timely diagnosis helps locate the malfunction once it happens; while predictive maintenance gives the alert before malfunctions occurs. Both of these mechanisms will help to improve the safety of the elevator and help pre-identify the faults which will reduce the maintenance expenditure.

(2) Security in preventive maintenance

Malfunction reporting and data uploading should be authenticated and encrypted. On-site engineers must also be authenticated, then authorized to access to maintain and repair.

Indoor Location Tracking

(1) Network bandwidth

High local network bandwidth (with reasonably low latency <100 MS) and provisions for backhaul to a local processing node with sufficient compute to allow the calculations necessary will be needed.

(2) Provisioning

The ability to add this capacity seamlessly within existing and future networks should be provided.

(3) Edge compute power

The local processing node should have a level of processing capability equating to, at time of publication, a server class machine. This can be achieved either by dedicated hardware or better via containerization within spare capacity on existing general purpose servers, or if future developments allow on an embedded compute node.

Lone Worker Safety

(1) Power efficiency

The footprint required for the location equipment is bulky due to the power/antenna requirements of current systems as well as the power requirements for maintaining GPS operation.

(2) Machine intelligence

False alarms can occur as the processing power of many edge devices does allow full

analysis to discern between normal movement, work activity, and an exceptional situation such as a fall. Nor is the processing sufficient to allow intelligent power management of the peripheral devices.

(3) Data transport cost

The cost of current data transport mechanisms is not justified by the level of data being transported.

Physical Access Control – Tailgating Detection/ Fire Detection via Surveillance Cameras

Besides requiring a greater local compute power, as described for the indoor location tracking, this use case would need additionally:

(1) Containerization

The ability to install the camera modules without hands-on access and with a high degree of modularity and security in existing environments is essential. In nearly all cases such applications will be add-ons to existing camera systems and fire/security infrastructures.

NEEDED CAPABILITIES

In order to ensure data privacy and prohibit any data or system tampering, IoT edge computing solutions are expected to be securely integrated with the cloud. Edge solutions also need to be centrally managed to minimize costs and to optimize lifecycle management across a wide range of edge devices. Data management and processing can take place at the cloud or edge, whichever approach is optimal for the specific scenario.

Required edge services include the following services:

- Persistence – to store IoT data on IoT gateways. IoT administrators can configure which data should be stored locally and set a data aging policy.
- Streaming – to analyze IoT data streams. IoT administrators can define conditions with adjustable time windows to identify patterns in the incoming IoT data as a basis for automated events. For example, the vibration, sound and other continuous data stream from a variety of sensors deployed in the machines, which can only be received in accordance with the sliding window order, should be analysed and compared to the existing rules just-in-time so as to detect the abnormal and initial subsequent transactions and notification of appropriate parties.
- Business transaction – to execute business transactions at the edge to provide continuity for critical business functions, even when the edge is disconnected from the cloud.
- Predictive analytics – to use predictive models for analyzing the IoT data. The predictive algorithm is constantly "being trained" and improved in the cloud based on all available data. The resulting predictive model is then sent to the edge and applied there.
- Machine learning – to apply machine learning algorithms at the edge specifically for image and video analysis.
- Visual analytics – to explore visually IoT data stored on IoT gateways. IoT data analysts can visually inspect the data collected at the edge. For example, after

an alert has been sent to the cloud, an analyst can dig into the details which led to the alert.

- 3rd party application hosting – to allow 3rd party application containers to be run on edge hardware, allowing decoupling between hardware and applications. For example an edge gateway might be used to run several services (camera, access control, AC management, elevators)
- End-to-end sophisticated management system –apply software-defined networking and other paradigms from 5G and other sources, to enable new business models on the edge intelligence based on tightly integrated services and networking.

CONCLUSIONS

The following conclusions can be drawn from the review and analysis undertaken in this article. The potential of edge intelligence in the Internet of Things, requirements, current technology gaps, and standardization need to be addressed to realize that potential:

- Edge intelligence (EI) is edge computing with machine learning capabilities. Data can be analyzed and decisions can be made by algorithms at the edge, i.e. very close to where the data is collected and where the machine and other equipment is controlled. This makes it possible to react autonomously (without a connection to the cloud) and with very short response times.
- Containerization will be important to deploy and manage edge intelligence consistently and economically. Containerization allows to encapsulate functionality, for example a machine learning algorithm, in a software package which can be deployed anywhere, e.g. in a public cloud, a private cloud, on premise, a micro data center on a shop floor, in a vehicle, within a 5G network or on an IoT gateway. This increases the efficiency of implementing new software development and allows optimizing the deployment according to customer specific requirements without additional programming effort. There currently exist no standards directly covering this technology, although there are many open-source initiatives, such as Docker and OCI.
- Common data models for edge computing node communication are essential to the success of edge intelligence. A common data model enables the interoperability between devices, communication protocols and software solutions from different vendors.
- Micro data centers will become more important in this process, for a number of reasons, including providing low latency and processing large volumes of data, thus avoiding transportation of such data to the cloud, which can be impracticable or costly.
- 5G networks will enable data centers at the edge and possibly industry-specific networks enabled by virtualization and software-defined networking principles. This allows customers to reduce costs and increase efficiency in a manner similar to the benefits provided by cloud

computing, as customers do not have to provision and maintain data centers at the edge.

- The best user interface is no user interface. Traditionally, user interfaces enable users to input data, analyze data and execute decisions. IoT makes manual data input largely obsolete, as data is collected automatically from sensors. Machine learning and artificial intelligence take over the data analysis and decision-making. Human interaction and interference are largely reduced.

The following recommendations are addressed to the various industries that will be impacted by the development and implementation of edge intelligence applications:

- Prepare for disruption of business and commercial models. During the last decade, we have experienced a change from the traditional software license model to the services model: software as a service, platform as a service, and infrastructure as a service. These services are typically located in the cloud, i.e. in centralized data centers. With the advent of edge computing, we will see an extension of these service models to the edge and combinations of traditional license and service models.
- Utilize 5G standards to facilitate edge computing and edge intelligence solutions. 5G networks have the potential to provide benefits of secure and scalable network connectivity.
- Agree on a common approach to orchestration and life-cycle management and to machine learning (tools, model implementation) to avoid market fragmentation. Their commoditization will drive down cost.

➢ Return to IIC Journal of Innovation landing page for more articles and past editions.