



**Unified software stack transforms
IIoT and AIoT efficiency**

Introduction

As IoT matures, there's a growing focus on edge intelligence as a means to minimise communication, eliminate latency and streamline processes and data analysis. Adding embedded intelligence at the edge is proving to be productive, secure and cost-effective across a range of verticals. Industrial organisations are stepping up their deployments of edge intelligence platforms so they can manage their deployed devices efficiently and reap the benefits of AI at the edge. Management platforms are essential enablers of edge and artificial intelligence and as hyperscale industrial IoT becomes a reality device design and development is being significantly impacted by edge capabilities.

Although it has taken much longer than expected, massive scale IoT is now here. Analyst firm Transforma Insights has reported that, at the end of 2023, there were 16.1 billion active IoT devices, a figure which it expects to grow to 39.9 billion in 2033, representing a compound annual growth rate (CAGR) of 10%. Annual device sales will grow from 4.1 billion in 2023 to 8.7 billion (a CAGR of 8%).¹

Even greater growth is expected in the industrial IoT (IIoT) sector which is forecast by Statista to experience substantial revenue growth from US\$238.40bn in 2024² to US\$454.90bn by 2029. That's a CAGR of 13.79% for the period. Worldwide, the IIoT market is witnessing rapid adoption of smart manufacturing technologies, revolutionising the way industries operate and boosting productivity. Part of this relies on the continued adoption of AI and the utilisation of edge technologies to control costs more tightly and to avoid latency caused by network delays and centralised cloud computing.

¹ <https://transformainsights.com/news/iot-connections-40-billion-2033>

² <https://www.statista.com/outlook/tmo/internet-of-things/industrial-iiot/worldwide>

AI empowers the edge

The adoption of AI edge processors is set to be driven by PCs over the next four years according to Omdia in the 2023 edition of its AI Processors for the Edge Forecast, which covers the market for artificial intelligence processors across ten device types including smartphones, PCs, tablets, smart speakers, robots and UAVs, among others. The firm projects the total market for AI processors at the edge is expected to grow from \$31bn in 2022 to \$60bn in 2028, a CAGR of 11%³, signalling greater adoption across a wide range of devices.

Part of this is due to a renewed focus on intelligent devices. The rise of cloud computing plus the availability of wireless connections has seen a trend towards deployment of cost effective, basic devices that simply gather data and send it to the cloud for processing. For non-urgent, low value use cases that rely on collection of a simple, constrained set of data points, this has been a practical approach but neither networking nor cloud are free resources and, as IoT use cases have become more complex, the volume of data involved, the complexity of the data and the need for it to be processed into usable insights for rapid response has accelerated.

All computing was done in cloud earlier which resulted in more costs and burden at the cloud and networking side, now this is being done on the edge device side to mitigate network demand and need for cloud.

The increase in analytics performed outside of cloud data centres is in line with Gartner analyst, Santhosh Rao's 2016 prediction that 'by 2025, 75% of enterprise-generated data will be created outside the traditional data centre or cloud.' It's now clearer that the 'outside' of Rao's vision is composed of these intelligent edge devices and that vast volumes of data will be created – and acted upon – by these devices. It won't be necessary for many routine tasks to send data to the cloud, freeing up network capacity for data that does need to be processed away from the device.

The increase in analytics performed outside of cloud data centres is in line with Gartner analyst, Santhosh Rao's 2016 prediction that 'by 2025, 75% of enterprise-generated data will be created outside the traditional data centre or cloud'



³ <https://omdia.tech.informa.com/pr/2024/feb/new-omdia-research-shows-ai-processors-for-the-edge-will-nearly-double-by-2028>

Too close to the edge?

Edge computing provides the processing power at the edge – either in the devices or in gateways – to enable data to be processed without having to transmit it to a centralised cloud and return the response over the network. This helps reduce costs in networking and cloud and can also reduce latency and simplify operations.

The vast volume of end devices therefore utilise different levels of intelligence to meet the needs of the use case they enable. Edge intelligence is seldom an either/or situation. Edge intelligence augments device intelligence and helps lighten the cloud computing and network load, enabling lower latency and reduced cost.

The edge therefore is variable in area with some intelligence very close to the edge in the form of on-device intelligence and other intelligence enabled much closer to an enterprise cloud function. Research firm IoT Analytics defines edge computing as: ‘intelligent computational resources located close to the source of data consumption or generation’. The firm defines edge as including all computational resources at or below the cell tower data centre and/or on-premises data centre, and there are three types of edges - thick, thin and micro.

- **Thick edge** describes computing resources, typically located within a data centre, that are equipped with components, such as high-end central or graphics processing units, designed to handle compute-intensive tasks/workloads such as data storage and analysis.
- **Thin edge** describes intelligent controllers, networking equipment and computers that aggregate data from sensors and devices generating the data.
- **Micro edge** describes the intelligent sensors and devices that generate the data.

Often fleets of devices such as soil sensors in agriculture will be able to gather data and transmit it to edge devices such as gateways where some processing can occur before some data is sent on to the cloud for further processing. In more complex use cases such as an AI-enabled colour sorter in a factory, the colour sorter will be able to separate

different colour products on an assembly line but also communicate data on performance and its own maintenance requirements to an edge device for more detailed processing.

With multiple types of edges defined, attention is focusing on edge AI, which describes the deployment of AI models on a device or piece of equipment at the edge, thus enabling AI inference and decision-making without reliance on continuous cloud connectivity. Artificial Intelligence IoT (AIoT) devices place significant demands on edge computing and helps to ensure maximised value can be derived from edge devices by enabling AI to be applied to them.

Segmentation of edge computing by category and type

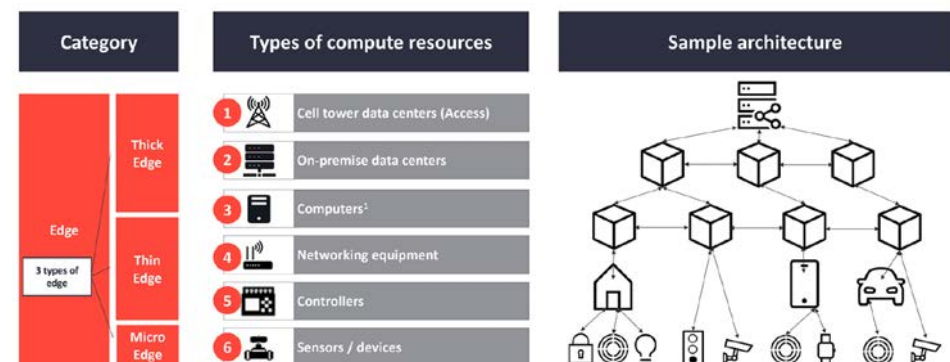


Figure 1: Three types of edges and commonly associated equipment

(source: IoT Analytics)


Million device management challenges

This is not a simple task and with huge volumes of devices, software and systems to manage AIoT, automation and edge intelligence are essential. Until recently, such software and platforms had heavy demands in terms of their consumption of network connectivity, computing resources and power. A new, lighter-weight approach is needed to provide a containerised OS that can manage the demands of AI at the edge within the constraints of typical IoT applications and use cases.

Eclipse Kanto, developed by the Eclipse Foundation, was initiated within Bosch IO to create the lightweight, connected, containerised and flexible system that the AI edge demands. Kanto is a modular, open source IoT edge system that enables devices for IoT with cloud connectivity, digital twins, local communication, container management and software updates - all of which are configurable and remotely manageable by an IoT cloud ecosystem chosen by the deploying organisation.

The bundle of enabling capabilities encapsulated within Kanto makes it possible to manage intelligence on edge devices so companies can get more value from their edge assets and processes and act on IoT data directly on the device as well as manage devices from the cloud. With effective management, device manufacturers can create new revenue streams with connected products and ensure agile development for hardware and software. Technology teams meanwhile can take advantage of Kanto's configurable and re-usable building blocks to connect the unconnected or build edge apps.

Eclipse Kanto, developed by the Eclipse Foundation, was initiated within Bosch IO to create the lightweight, connected, containerised and flexible system that the AI edge demands



Designed for complex IoT devices

Kanto has been specifically developed to meet the needs of complex IoT devices that have constrained hardware resources, near real-time requirements, different device software, heterogeneous data sources and the ability to operate without connections. This makes it well-suited for deployments such as software-defined vehicles, connected machines and connected manufacturing, smart appliances and smart building use cases.

Based on lightweight native components, open protocols and standard containers that ensure availability on hardware with minimised integration work, reduced complexity and scalable applications, Kanto is a modular IoT edge software that enables devices for IoT with all the essentials such as cloud connectivity, digital twins, local communication, container management and software updates - all configurable and remotely manageable by an IoT cloud ecosystem of choice.

Kanto has four key benefits:

- **Lightweight** – Kanto is deployable on various constrained edge devices with different hardware capabilities so it can scale hardware and connect edge devices and data sources simply.
- **Containerised** – Kanto helps to develop edge applications using the most appropriate languages and frameworks and assists integration with a wide range of device software and cloud-native technology in a uniform way.
- **Connected** – Kanto connects to the most-fit-for-purpose public IoT platforms, can help manage devices and device data from the cloud with minimum customisation, and enables rapid prototyping for new products and solutions.
- **Flexibility** – Users can choose and combine re-usable, configurable and interoperable building blocks to bring the right IoT edge capabilities to their deployment depending on the hardware and use case.

OEMs and companies need to continue digitalisation and to build connected products and digital services. To maximise edge capabilities, minimise costs and achieve fast time to market, software such as Kanto can help reduce the complexity of edge software enabling user organisations to focus on their core business opportunities, differentiation and fast innovation. Kanto can help speed up development, supporting fast prototyping, and the rapid migration of cloud-native applications to edge devices.



Open standards, common APIs

An important aspect of Kanto is its support for open standards and common APIs which can be bundled in re-usable, configurable building blocks. This ensures interoperability between different apps and ease of use. The connectivity, via message queueing telemetry transport (MQTT), to the cloud helps ensure there is no vendor lock-in. Kanto can support other Eclipse IoT projects, the Bosch IoT Suite and accommodates adapters for Azure and Amazon Web Services. It provides the device management platform, containerisation, maintenance and orchestration for different operating systems and containerisation technologies.

The open source ecosystem is further enriched with Kanto in the middle to ensure ease of operation and support for massive scale. Kanto has been utilised internally within Bosch for several years and the next step is to take the mature and robust technology, which incorporates knowledge and insights from edge experts with more than 20 years of experience in both edge and IoT, and use it to establish partnerships for Kanto with semiconductor companies and others.



Why Kanto?

The Kanto software stack has compelling advantages for the management of the ever-growing base of intelligent edge devices because it has no licence fees and no vendor lock-in. It's therefore ideally placed to help semiconductor companies achieve faster time-to-market and to bring companies' IoT solutions to market rapidly. Kanto is also positioned to help IoT organisations get fragmentation of systems under control. By acting as a connection layer, Kanto can be used by an OEM that wants to develop a new IoT device that works with the systems it needs to use and offers all the features at a click of a button. Kanto is being used by one large Asian manufacturing company, for example, that fits engines to very large vehicles. This involves very expensive components being deployed for global markets and the hardware is installed on top of the engine to monitor efficiency. Ecosystem partners can connect confidently with common hardware and enable on-boarding of applications. Being open source based only increases the applicability and value of the ecosystem.

Kanto also brings with it the capability to perform firmware updates and corresponding cloud services that can help abstract away operational complexity and streamline development of features and services.

Being open source based only increases the applicability and value of the ecosystem

