Bringing Intelligence to the Edge:

Exploring Edge Computing

Edge Computing: The old and the new

Edge Computing is not a new concept. As a subset of distributed computing, it has been talked about for many years as a mechanism to move latency sensitive applications closer to their source of data.

However, over recent years, several factors have combined to really drive the adoption in this space moving Edge out of the hype cycle and into a modern platform choice to sit alongside its counterparts in cloud and the data center. These factors include the rapid increase in the capability of networking technologies such as 5G, the rise of data derived from IoT, the AI revolution and the power of microservice based applications. Further, with the ever-increasing number of available device types, Edge Computing is further converging the worlds of Information Technology (IT) and Operational Technology (OT).

Customers in many vertical industries are not only using Edge solutions, they are also dependent on them to run a data driven business that can make real time changes. As customers start driving their technological reach out to the Edge, strategic consideration is required as deploying solutions successfully can be both nuanced and bespoke.

We know Edge is now a central pillar of a hybrid platform strategy, and the build out of further distributed applications to drive AI and other initiatives will be part of many organisations' strategies. What will keep CTOs awake at night is how to ensure the solution efficiently and securely delivers its expected business value as Edge devices, Edge networks and local data storage proliferate in the coming years.

Customers in many vertical industries are not only using Edge solutions, they are also dependent on them to run a data-driven business that can make real-time changes.

The Fourth Industrial Revolution

Industry 4.0 is a term to describe the integration of intelligent digital technologies into industrial processes by encompassing Edge-based technologies such as IoT networks, AI, Big Data, robotics, and automation. We are already seeing how the implementation of these solutions are revolutionising vertical industries with uses cases such as:

Public Sector

Providing secure border control by processing biometric data such as facial recognition via local cameras and sensors at border checkpoints.

Smart Cities

Managing urban infrastructure by optimising traffic lights, waste management, and energy grids.

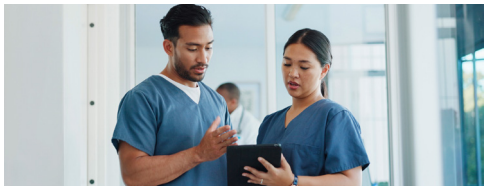

Healthcare

Edge devices such as wearables and medical sensors securely collect patient data and transmit it to healthcare providers.

Defence

Soldiers can be digitally enabled with devices such as wearable sensors, smart helmets, and communication gear, collecting data from various sources, including GPS, biometrics, environmental sensors, and weapon systems.

Retail

Through the deployment of sensors and cameras in stores, real-time data can be gathered on customer behaviour, such as foot traffic patterns, dwell times, and product interactions.

Finance

Insurers can assess risk and make underwriting decisions in real-time by monitoring data from IoT sensors in homes. Understanding risks such as water leaks or fire hazards enables insurers to price services more accurately.

Finding the Edge: the 6 core areas to consider

The use cases for Edge, clearly aligned to business outcomes, will affect target architectures. Once these have been created, successful Edge adoption requires the knowledge and service adjacency across several key areas to successfully deploy, maintain and evolve solutions.

Being able to bring infrastructure, applications, data, security, and networking together to meet the specific requirements of the use cases, in addition to any operating model changes required, is key to reducing the risk and accelerating the delivery of Edge solutions.

1. Data

Data at the Edge is no longer bound within the confines of a secure data center or cloud platform which creates both opportunity and challenge. AI is a key use case for Edge, and Edge is a key enabler for AI to meet some of the many use cases. As AI is a data-heavy, compute-intensive technology it becomes a perfect candidate for Edge as models and processing can be driven at source with minimal network latency. This is one of the reasons Edge is driving data creation at an unprecedented scale, so much so that it's predicted that:

- By 2025, more than 50% of enterprise-managed data will be created and processed outside the data center or Cloud.[1]
- By 2027, 20% of large Enterprises will have deployed an Edge management and orchestration solution, compared to less than 1% in 2023.[1]
- The number of IoT devices will triple from 2020 to 2030, and the data they produce will grow even faster.[1]

Edge Data Management is evolving as a focussed discipline for storing, managing, and processing data from the nearest source to reduce latency, mitigate short delays and enable high throughput. The application and use case will determine if data is sent back to a centralised server or cloud platform or is processed locally. Clearly, the further the data travels the more consideration needs to be given to security and factors such as encryption in transit, but other key considerations are:

- Backup and recovery
- Performance and reliability
- Data storage both locally and remotely
- Data partitioning where huge amounts of data are being collected (i.e. from sensors)
- How to use AI services on Edge data

1. 2023 Gartner IT Infrastructure, Operations & Cloud Strategies Conference

2. Infrastructure

Today there is a proliferation of Edge in various form factors. They may be IoT sensors, cameras, mobile devices, drones, and as many other connected devices as you can imagine from the Internet of Things. The hyperscalers now provide localised 'zones' to run cloud services in metro areas or supply hardware devices that can run certain services decoupled from the network. These include AWS Outposts, AzureStack HCI and Google Distributed Cloud.

This is where we bring in the more granular concept of what Gartner have termed 'Cloud-Out versus Edge-In'. Edge-in is about building out CSP independent solutions and connecting them back to the Cloud. Such as devices in stores or wearables. Cloud-Out is about extending what you already do in Cloud to the Edge. Such as a HCI device from a hyperscaler.

The challenges of managing these devices, outside of the traditional perimeters of the data cener or cloud platform is twofold. Firstly, how do you perform standard operational services on these devices, such as monitoring, alerting, patching and support? Secondly, a different type of management service is required. At the Edge, you could have thousands of locations and all sorts of diversity in the field. One location will be cold, the other one will be hot. One location will have a mobile cellular network available another may be Wi-Fi.

As reliability is crucial for Edge applications, the infrastructure must be flexible and efficient emphasising automation and visibility in management platforms to simplify decision-making. By failing to consider these aspects organisations are in danger of repeating the mistakes made 10-15 years ago when IoT struggled to gain traction due to a lack of standardisation and support in device management.

Mapping Edge-In and Cloud-Out Services and Locations

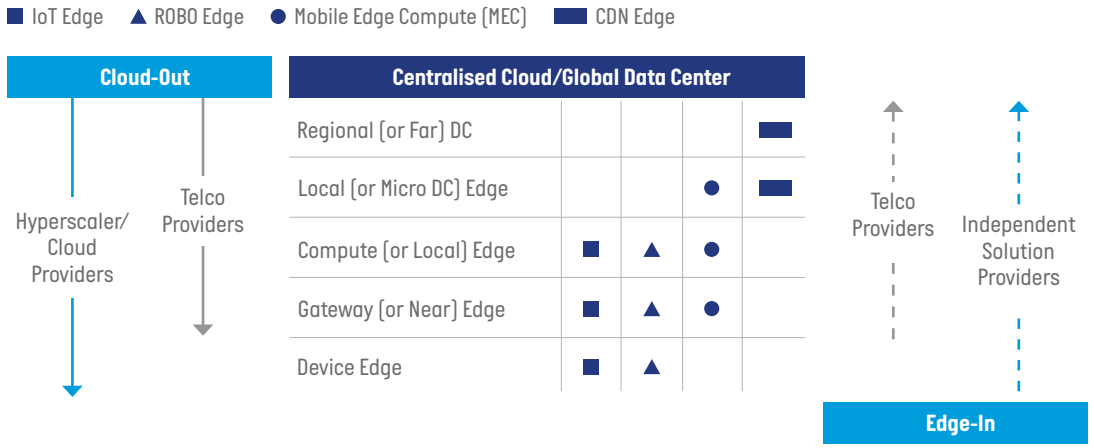

Figure 3: Mapping Edge-In and Cloud-Out Services and Locations, gartner.com

3. Applications

To successfully meet the use cases required applications must be deployed to run at the Edge and be executable across various device types. These applications need to be built with key requirements in mind around low latency, local data access and eventual consistency with a centralised platform. APIs are used to abstract the underlying edge infrastructure allowing an application to interact with the edge infrastructure without needing to understand the underlying device technology in depth. The APIs at the edge can also support applications working offline (asynchronously) and working synchronously to connect back to a central platform (such as the cloud).

Microservices and Container solutions have emerged as a common method of running services or executing software in smaller form devices and utilising these API layers. Containers are an especially good fit for edge applications providing modularity, segregation, and immutability. Applications can be deployed on many different edge tiers, each with their own unique resource characteristics with the containers representing them able to scale up or down depending on changing resources and other conditions.

This is further supported by orchestration tools, like those provided by Kubernetes, that can recognise edge locations as they would data centre or cloud platforms and being able to manage them from a centralised container control plane, such as RedHat OpenShift or the related services delivered by hyperscalers.

4. Networking

There are two things that will shape the network requirements for your Edge Computing solution;

- Location
- Application

These are derived from use cases, combined with a data flow mapping exercise. As examples:

If the use case means locations are remote, this will inform available connectivity options, and if the application demands a certain level of bandwidth the options will be refined. Options for connecting to a remote location may be reduced to Public 5G, LTE or Low Earth Orbit Satellite Internet and then the bandwidth demands of the application and subsequent data transfer may require SD-WAN to maximise performance. Conversely the use case may be defined by a lack of stable bandwidth to a location and so Edge Computing is the solution, but it will still have its own compute, gateway, and device edge connectivity requirements.

The location may be small so minimising the footprint of the Edge solution means using SmartNICs to offload network functions to reduce load on the Compute platform or even to deploy a software version of a hardware network appliance. The location environment may be harsh either in terms of exposure to weather, or permanent high or low temperatures, which will require specialised equipment to solve. The use case may also dictate the type of network technology required, for example the connectivity for a large factory floor or warehouse with a high density of connected devices may be better provided by Private 5G rather than traditional Wi-Fi.

Finally, the distribution of Edge locations and their requirements to connect back to the mothership, particularly if that is Cloud, may be better served by a modern connectivity platform that uses an Internet on ramp to a private, high-speed backbone, that has proximity to the hyperscalers. Understanding the use case and its location and application requirements is fundamental to how you will architect your network.

5. Security

Security good practices will apply to Edge as much as they do to any other part of the business, from the user to infrastructure and networking, through to shifting left towards the application, but there are some specific considerations. For example:

- Zero Trust principles should be applied to Edge. Identity driven least privilege is critically important as it is anywhere else.
- Secure Access Service Edge (SASE), particularly through Zero Trust Network Access (ZTNA), will ensure user access to Edge workloads is secured and reduce the attack surface of those workloads.
- Securing IoT/OT edge devices will require dynamic segmentation both from each other and other parts of the network.
- If a location is remote an extra focus on physical security may be required as well as encryption at rest and in flight.
- SmartNIC technology comes into play again to minimise the effects that security controls, such as firewalls and IDS/IPS, have on the footprint of the solution.

The challenge is to provide consistent policy and enforcement while dealing with the nuances of the use cases and diversity of the components of Edge solutions.

6. The importance of platform thinking

Although we have talked about technology domains individually here, we are not saying that they need to be architected individually. A focus on maximising the availability of modern, consolidated, multi-domain platforms will be key to Edge success, as this will accelerate deployment, streamline operations, and ultimately provide a stronger security posture.

Platform engineering will be a requirement to address any vendor platform gaps to ensure maximum efficiencies in the final solution and ultimately maximise the business value of Edge.

How Computacenter can help

As more of our customers begin to drive their technology to the Edge, we have a defined set of service areas designed to deliver the assessment, architecture, management, and transformation our customers need. Computacenter have dedicated business units and specialists across each of the core consideration areas, therefore placing us in a unique position to help our customers.

Computacenter's **Advisory & Assessment** Services can advise you on the best technology choices, designing the solution end-to-end, including integration into existing environments and the Cloud. We can build workload optimised platforms for artificial intelligence, databases and IoT solutions.

Using the breadth of our portfolio expertise and technology partners, we **Engineer & Implement** the most appropriate Edge platforms for our customers to achieve their digital and transformational requirements.

And finally, to deliver operational efficiencies, we offer services to help **Adopt & Operate** platforms. From training and day 2 support through to ongoing optimisation, scalability planning and lifecycle management, we can help customers to get the most value from their Edge solutions.

Let Computacenter take you on your journey to the Edge!

INFRASTRUCTURE SERVICES

DC & EDGE PLATFORMS

RELOCATION & MIGRATION

OPERATIONAL RESILIENCE

Platform

Data

Network

Security

Advise & Assess

- Requirements gathering
- Business case creation
- Platform recommendations/ Battle Days
- Architecture design
- Data assessment and classification

Engineer & Implement

- Deployment services
- Integration to existing systems
- Networking and connectivity
- Security implementation
- Compliance and regulation

Adopt & Operate

- Training and support
- Scalability planning
- Performance optimisation
- Lifecycle management
- Circular Services

The Office of the CTO

Computacenter's Office of the CTO (OCTO) team leads in the exploration and application of technology products and delivery methodology to aid the digital transformation of our customers. As a team of cross-functional technologists with extensive industry and IT experience we deliver thought leadership, advice, and real-world implementation experience to help our customers achieve their goals.