LF NETWORKING

LF EDGE

Edge Networking: An Introduction

An exploration of the edge opportunity, the role of open source and Linux Foundation projects, and how to participate

Please direct any questions to **lfn-info@linuxfoundation.org**.

1 What is the Edge?

"The edge" represents a new computing paradigm in which compute and storage are located at the edge of the network, as close as possible to the physical location where data is generated and collected.

The network edge itself is still being defined by industry groups and there are still some fluid areas. One useful distinction however to point out is the **infrastructure edge** such as a micro-modular data center or carrier cell-site and the **device edge** such as a smartphone, drone, intelligent video camera or connected vehicle. Edge compute functions can be hosted on either side, operating as an extension of the centralized cloud in the context of an overall application or workflow that's managed in the cloud but executed at the edge.

1.1 What is driving edge deployment today?

Both enterprise and consumer applications benefit from the real-time decision-making, reduced connectivity costs and improved security associated with edge compute. Some examples include:

- **Augmented Reality & Virtual Reality (AR/VR)** which, beyond consumer applications like Pokémon-type games and immersive sporting experiences, brings business efficiencies to manufacturing and maintenance operations. Edge-hosted AR/VR eliminates the dizziness caused by excessive latency and slow refresh rates, while moving compute into edge servers reduces the need for high processing power on AR goggles.
- **Video surveillance**, whether for security purposes or as part of operations in manufacturing, agriculture or logistics, benefits from video data being processed in the camera itself or on an edge server. Processing data at the edge accelerates time- and safety-critical functions such as facial recognition or machine alignment detection, while only selected events need to be uploaded to the cloud, reducing backhaul bandwidth.
- **Cloud gaming** requires ultra-low latency that is consistent for all players, for Massively Multiplayer Games (MMPGs) to provide a fully-immersive experience. Edge compute ensures this latency, while also reducing device costs by shifting game processing to edge servers.
- **IoT** leverages edge compute both to reduce bandwidth costs through processing data locally and to improve agility through real-time on-premise decision-

making. Where cloud connectivity is intermittent, edge compute enables processes to function autonomously. Includes Industrual IoT (IIOT).

- Other applications in areas like **healthcare, manufacturing, smart cities, smart buildings, autonomous vehicles** and **smart retail** also leverage edge compute to transform existing processes, create new user experiences and enable new business opportunities.

Most of these applications will use Artificial Intelligence (AI) and Machine Learning (ML) to enable autonomous operations and improve the user experience. For example, a smart security camera will use AI-based vision processing to distinguish authorized employees in a commercial space, only sending an alert to the cloud if a person doesn't match pre-established parameters.

Companies of many types are pursuing business opportunities associated with delivering edge services, including:

- **Communications Service Providers (CSPs)** own and operate both fixed and mobile infrastructure that provides cloud connectivity as well as physical Points of Presence (PoPs).
- **Colocation Hosting Providers** provide Multi-Tenant Data Centers (MTDCs) at the edge.
- **Cloud Service Providers** are extending their public and hybrid cloud solutions to encompass the edge.
- **Edge Data Center Providers** are deploying collocated, multi-tenant micro-modular data centers at the edge in innovative form-factors like shipping containers.
- **Content Delivery Network (CDN) Providers** are extending their worldwide network of edge locations to host third-party edge applications beyond CDN.

1.2 How is edge compute transforming networks?

Many edge compute applications work perfectly without the need for high-speed networks. Low-bandwidth IoT use cases are deployed using low-power, wide area networking protocols like Long Range Wide Area Network (LoRaWAN), Bluetooth Low Energy (BLE) and Narrowband IoT (NB-IoT). Likewise, WiFi and 4G/LTE are adequate for many video surveillance, AR and online gaming applications.

With the advent of 5G, however, businesses will launch new types of edge applications that were not possible with previous generations. Some of the benefits of 5G for edge compute include:

- **High connection density** of one million devices per square kilometer, supporting the proliferation of IoT sensors and wearables.
- **High bandwidth**, with data rates as high as 10 to 20Gbps enabling new use cases in streaming video, autonomous vehicles and haptic applications.
- **Ultra Low latency** for edge applications like industrial robots and drone navigation.
- **Network slicing** that allows an operator to share infra and radio spectrum for efficient utilization of infrastructure and do things such as deploying only the functions required to support a specific customer, for greater operational efficiency than providing full functionality to devices that don't need it.

Other benefits include involving the network edge in other aspects of 5G development including automation, orchestration, AI/ML, and cloud native development approaches. 5G also enables new business models and use cases, and ushers in a new era of innovation to the industry. As more companies deploy and operate their own private networks, whether based on 4G/LTE or 5G, an increasing number of enterprise-focused edge services such as factory robotics, industrial AR, smart retail and smart buildings will benefit from deterministic performance characteristics. Similarly, data security will be enhanced by ensuring that sensitive information is retained on-premise rather than being transferred over a public network.

2 Open Source and the Edge

2.1 Open source is indispensable to edge

Open source has become the de facto approach for creating software. Innovative companies today are evolving away from proprietary, fixed function, closed infrastructure in order to open up new open source business models that create operational value, generate new revenue streams, reduce costs and enhance competitiveness. As companies and developers use open source code to build their own commercial products and services, they also gain strategic value through contributing back to those projects. Open source reduces time and optimizes effort, while improving efficiencies and interoperability.

Companies have recognized that the proliferation of edge applications across a growing number of market segments causes an increase in the complexity of the overall edge ecosystem. Besides data center providers, this ecosystem includes suppliers of servers, virtualization platforms, orchestrators, management systems, analytics and of course the applications themselves. This complexity is driving the adoption of open source software and standards at the edge.

Rob High, VP and CTO of IBM Watson, summarized the importance of open source to standardization within the edge industry: "…. the No. 1 thing you begin to realize is that this industry is at risk of imploding on itself if it does not solve the problem of creating a standard way of managing it, [and] creating a set of standards that developer communities can begin to form and create ecosystems from…. We think that the only way to achieve that standardization in a meaningful way is through open source."

2.2 What role does open source play?

Open source projects have emerged to address needs and business opportunities throughout the end-to-end architecture and lifecycle of edge services. The community is actively working on projects such as:

- **Edge infrastructure platforms**, capable of hosting edge applications and network functions that are virtualized and/or containerized.
- **Orchestration tools** that automatically provision the infrastructure services required to efficiently deploy applications and network functions at scale, then remove them when no longer needed.

- **Management tools** that automate the processes associated with keeping the infrastructure, the network functions and the applications secure, up-to-date and available.
- **Monitoring tools** that provide real-time analytics and telemetry used by the management and orchestration subsystems in order to diagnose and resolve problems throughout the system, while optimizing resource allocation and performance.
- **Network functions**, running in Virtual Machines (VMs) or containers, that implement the underlying network infrastructure at the device edge, enterprise edge and/or telco edge.
- **Applications**, running in VMs or containers, that implement user-visible or device-visible edge services, such as AR, machine vision, robotic control, drone navigation, gaming, video analytics etc.
- **Security** at the edge presents unique challenges of securing and managing compute devices located outside the perimeter of a company's corporate networks and data centers and involves things like device identity and onboarding.

2.3 Interdependence of edge, 5G and cloud-native transformations

In order for edge applications such as IIoT to realize their full business potential, they must be implemented and operated within an environment that allows for highly automated deployment at large scale (for example millions of sensors in a manufacturing plant), efficient updates on a frequent basis, the elimination of unplanned system downtime and hardware platform independence. These objectives are best realized by the application of cloud native principles.

Cloud native applications utilize modern Development and Operations (DevOps) methodologies that enable the continuous delivery of loosely-coupled, platform-independent microservices. The microservices architecture ensures automatic scalability that eliminates downtime due to errors, as well as rapid recovery in the event of application or infrastructure failures. A cloud native edge provides rich, real-time insights into operations through telemetry, used by predictive analytics for ongoing optimizations of operational processes.

As companies leverage 5G networks to deploy edge applications that scale to high connection densities and use network slicing to deliver customized services to large numbers of individual endpoint devices, cloud native architectures will be the optimum approach to ensure the scale, automation, reliability and security that is required.

3 Linux Foundation: Bringing Innovation and Community to the Edge

The Linux Foundation (LF) is a non-profit technology consortium founded in 2000 to standardize Linux, support its growth and promote its commercial adoption. LF and its projects have more than 1,500 corporate members from over 40 countries. LF also benefits from over 30,000 individual contributors supporting more than 200 open source projects.

3.1 Linux Foundation Networking (LFN) capabilities for the edge

As part of the LF organization, Linux Foundation Networking (LFN) comprises the world's largest community focused on networking technologies. LFN brings together leading projects in Software Defined Networking (SDN), Network Functions Virtualization (NFV), Management and Orchestration (MANO), automation, analytics, dataplane acceleration and more.

LFN projects are deeply integrated with edge compute initiatives, including strong collaboration with the Linux Foundation Edge (LFE) community (details below) to add networking-focused capabilities and features to edge blueprints and use cases.

The **Open Networking Automation Platform (ONAP)** provides an example of support for edge use cases within an LFN project. ONAP is a platform for real-time, policy-driven orchestration and automation of both physical and virtual network functions. Through collaboration between the LFN and LFE communities, ONAP is being extended to provide support for onboarding edge clouds from distributed edge locations.

In another example of collaboration between LFN and LFE, a reference platform for edge cloud use cases is being developed within the **Open Platform for NFV (OPNFV)** project. The project is focused on a defined set of edge compute scenarios, addressing edge-specific constraints such as the limited space and power resources, unattended operation, scale up to tens of thousands of locations, network infrastructure acceleration and heterogeneous platforms.

A number of LFN projects are leveraged directly by the LFE community for edge use cases. For example, both **OpenDaylight (ODL)**, a modular, open platform for

customizing and automating networks, and **Tungsten Fabric**, a virtualization solution for network connectivity and security, have been used by LFE as SDN infrastructure for network edge projects based on the **Akraino** edge stack (details below).

3.2 Linux Foundation Edge (LFE) overview

As a peer to LFN, Linux Foundation Edge (LFE) was founded in 2019 as an umbrella organization to establish an open, interoperable framework for edge computing independent of hardware, silicon, cloud or operating system.

LFE comprises the following open source projects:

- **Akraino** is a software stack that supports high-availability cloud services optimized for edge computing systems and applications. It offers users new levels of flexibility to scale edge cloud services quickly, to maximize the applications and functions supported at the edge and to help ensure the reliability of systems that must be completely functional at all times. Release 2, available since January 2020, includes 6 blueprint families and 14 blueprints. These blueprints support use cases in both telco cloud, such as **Radio Edge Cloud (REC)**, and network cloud, such as **Unicycle**, which provides a framework for defining and managing infrastructure lifecycle.
- **Baetyl** (pronounced "Beetle") is a general-purpose platform for edge computing that manipulates different types of hardware facilities and device capabilities into a standardized container runtime environment and API, enabling efficient management of application, service and data flow through a remote console both in the cloud and on premise.
- **EdgeX Foundry** is a vendor-neutral, loosely-coupled microservices framework that provides the choice to plug and play from a growing ecosystem of available third-party offerings or to include proprietary innovations.
- **Fledge** is a framework for the industrial edge focused on critical operations, predictive maintenance, situational awareness and safety. Fledge is architected to integrate IIoT, sensors and modern machines with the cloud as well as with existing legacy systems.
- **Home Edge** is a robust, reliable and intelligent home edge computing open source framework, platform and ecosystem. It provides an interoperable, flexible and scalable edge computing services platform with APIs that can also be used with libraries and runtimes.
- **Open Glossary of Edge Computing** provides a concise collection of terms related to the field of edge computing, improving communication and

accelerating innovation through a shared vocabulary, offering a vendor-neutral platform with which to discuss compelling solutions offered by edge computing.

- **Project EVE** is an edge computing engine that enables the development, orchestration and security of cloud native and legacy applications on distributed edge compute nodes. Supporting containers, clusters VMs and unikernels, it provides a flexible foundation for IoT edge deployments.

In addition to the baseline project descriptions above, a number of technologies important for edge use cases are being investigated and incubated within these projects. Within the Akraino community, for example, experts are exploring options for cloud native offloading stacks in response to the significant challenges associated with supporting 25G, 50G and 100G links on general-purpose processor cores.

3.3 Other LF initiatives

Outside the LFN and LFE organizations, other communities and projects are relevant to edge applications, including:

- **Cloud Native Computing Foundation (CNCF)**, an LF project, promotes the adoption
of cloud-native computing, aiming to establish a vendor-agnostic community of developers, end users and IT technology and service providers to collaborate on technologies such as containers, microservices and service mesh, with the goal of developing and deploy scalable applications on cloud computing platforms.
- CNCF established the **Telecom User Group** for CSPs and their vendors who plan to use cloud-native technologies in their networks, in order to collaborate on requirements, best practices, gap analyses and similar documents.
- **Common NFVi Telco Taskforce (CNTT)**, was created in partnership by GSMA and LF in 2019 to operate as an open committee responsible for creating and documenting a Common NFVi Framework. Co-hosted by GSMA and OPNFV, the goal of the framework is to create reference implementations and architectures to reduce cost, time-to-market and complexity of telco operations in development and adoption of VNFs and CNFs. **A CNTT Edge working group** has been formed as well.

4 How to Get Involved

4.1 LF Networking

Whether your organization is a networking solutions vendor, system integrator, end user network operator, cloud services provider or an enterprise, joining the LFN community connects you with the innovative technical projects, companies and developer communities that are transforming the networking industry. Membership showcases your support for community-driven, open source solutions, while accelerating the development and adoption of open source SDN, NFV, orchestration, automation, and more.

The **Join** page on the LFN website provides information on joining LFN as a member, explaining the processes for both existing LF members and non-members. See **this Getting Started page** to learn about the many ways to participate in LFN. There's also a link to an **Inquiry** page where interested parties can ask specific questions and obtain additional information.

4.2 LF Edge

Industry-leading organizations across telecom, cloud and enterprise technology have joined LFE to support the development of open, interoperable frameworks for edge computing independent of hardware, silicon, cloud or operating system. Members are collaborating across company and geographic boundaries to grow and shape the ecosystem for edge and IoT technologies, use cases and applications.

The **Join** Join page on the LFE website provides information on joining LFE as a member, explaining the processes for both existing LF members and non-members. **The LFE Getting Started Guide** provides information for getting involved. There's also a link to an **Inquiry** page where interested parties can ask specific questions and obtain additional information.

Please direct any questions to **lfn-info@linuxfoundation.org**.