EdgeCloud: State of the art 2020

Henrik Abrahamsson, RISE Justin Lex-Hammarskjöld, KTH,
Ian Marsh, Nicolae Paladi, RISE AB

1 The Edge Cloud ecosystem

Edge cloud is a diverse area which is under rapid development. The landscape is dynamic and fragmented. There are many different organizations and companies (and alliances of these) that drive development forward, come up with new initiatives, and present new technology and their own solutions. Edge solutions are specified and promoted by many different companies, open source projects, standardisation organisations and industry groups. In this short-paper we present an overview of ongoing initiatives, standardisation work, and state-of-the art. A history of containers can be found in [Osnat2020].

1.1 Edge cloud features

Edge cloud provides distributed computing and storage at the edge of the network. It offloads end devices by having data and processing in the network. It gives lower latency for applications compared to centralized data centers. It also gives less network traffic (since data is kept at the edge) which saves network resources. There are many ways to accomplish distributed computing and storage, and edge computers can be anything from a microcontroller to an entire room full of servers. Edge computing infrastructure is broadly defined, and can be summarized as "everything that's not in a central datacenter" [Reznik 2018]. Edge is a location. Edge cloud is a concept, not a single technology or product. Edge cloud involves many different possible techniques and is often presented bundled with (or as an enabler for) low-latency use cases such as: connected vehicles, IoT, smart cities, robotics, next-generation industries, gaming and virtual reality.

Many cloud providers are extending their services towards the edge. For instance, Amazon Web Services (AWS) provides an edge solution in AWS IoT Greengrass. It is an edge run-time and cloud service where devices can act locally on the data they produce. Similarly, Microsoft provides Azure IoT Edge - *"Cloud intelligence deployed locally on IoT edge devices".* It uses containers to move workloads to the edge for lower latency. Microsoft also provides EdgeML where training of machine learning algorithms can be done in the cloud, but the trained models are run in the edge or endpoint IoT devices (and that's where the classifications or predictions are done). Another angle, Edge computing from a communication service provider perspective, is outlined in the whitepaper [Ericsson GFMC-20:000097] "Edge computing and deployment strategies for communication service providers", published in February, 2020. Some perspectives on Edge cloud in 5G and cloudification of access is also given in the book 5G Mobile Networks: A Systems Approach.

The cloud provider and network operator perspectives on edge clouds are the two main paths forward. But there are also other possibilities. For instance, from a network research perspective Peterson et al [Peterson2019] argue for a "democratization of the network edge": *"On the one hand, cloud providers believe that by saturating metro areas with edge clusters and abstracting away the access network, they can build an edge presence with low enough latency and high enough bandwidth to serve the next generation of edge applications. In this scenario,* *the access network remains a dumb bit-pipe, allowing cloud providers to excel at what they do best: run scalable cloud services on commodity hardware. On the other hand, network operators believe that by building the next generation access network using cloud technology, they will be able to co-locate edge applications in the access network. This scenario comes with built-in advantages: an existing and widely distributed physical footprint, existing operational support, and native support for both mobility and guaranteed service. While acknowledging both of these possibilities, there is a third outcome that not only merits consideration, but is also worth actively working towards: the democratization of the network edge."* Peterson et al argue that the softwarization and virtualization of the access network opens the door for anyone (from smart cities to manufacturing plants) to establish an access-edge cloud and connect it to the public Internet. And it should become *"as easy to do this as it is today to deploy a WiFi router"*.

1.2 Academic work and research papers on Edge computing

A vast number of research articles have been published on Edge computing and related paradigms. Pham et al., "*A Survey of Multi-Access Edge Computing in 5G and Beyond: Fundamentals, Technology Integration, and State-of-the-Art*" [Pham2020] lists 329 related papers. Mach and Becvar, "*Mobile Edge Computing: A Survey on Architecture and Computation Offloading*," [Mach2017] lists 124 papers. Abbas et al, "*Mobile Edge Computing: A Survey,*" [Abbas2018] cites 312 papers, and Yousefpour et al, "*All one needs to know about fog computing and related edge computing paradigms: A complete survey.*" [Yousefpour2019] lists 450 related papers. There are many Edge-related papers published every year in conferences like IEEE Infocom. There are also specialized workshops like ACM EdgeSys on the topic.

2 Edge initiatives: standards, open source projects and industry alliances

2.1 Standardisation organisations

2.1.1 European Telecommunications Standards Institute (ETSI) and 3GPP

The European Telecommunications Standards Institute (ETSI) initiated standardisation work for mobile edge computing already in 2014. ETSI ISG MEC and 3GPP have both worked on their own architectures for edge computing within their scopes. Their common purpose is to create an open and standardized IT service environment for hosting and supporting third-party applications in edge environments. 3GPP SA6 is specifying the 3GPP UE and network's relationship with the EDGE APP system. 3GPP SA2 is contributing to edge computing for 5G, with a main focus on DNS optimization for optimal discovery of/routing to edge clouds. MEC Platform/Edge Enabler Server and the MEC Applications/Edge Application Servers are at the center of ETSI ISG MEC and 3GPP SA6 architectures. The ETSI White Paper [ETSI36] *"Harmonizing standards for edge computing - A synergized architecture leveraging ETSI ISG MEC and 3GPP specifications"* (July 2020) describes how the ETSI and 3GPP standards may be combined when it comes to deployments.

2.2 Open source projects

2.2.1 LF Edge

LF Edge is an umbrella organization that started in January 2019. The objective is to establish open and interoperable frameworks for edge computing independent of hardware, operating system and cloud. LF Edge has 77 member companies (at the time of writing) and is supported by the Linux Foundation. A number of projects are conducted under the LF Edge umbrella including: Akriano, EdgeXFoundry, Project EVE, *Home Edge Project, and State of the Edge.*

Akriano and EdgeXFoundry were existing Linux Foundation projects before the LF Edge umbrella was initiated. Akriano is an open source edge software stack for carrier, provider, and IoT networks. Akriano provides a set of blueprints. A blueprint is a tested and approved declarative configuration of the entire stack including edge platform, API and applications. Akraino was initiated by AT&T and Intel. AT&T provided seed code to enable carrier-scale edge computing applications to run in virtual machines and containers. Akraino uses many other open source components within its blueprints including Openstack and StarlingX.

EdgeXFoundry is an open-source platform for industrial IoT edge computing. It provides a set of micro-services for data transfer between sensors and edge cloud applications. Project EVE focuses on IoT edge virtualisation, while the *Home Edge Project* provides a home edge computing framework where TVs, fridges, and washing machines are considered Home Edge devices and are monitored by an Home Edge Orchestrator.

The State of Edge project produces a yearly market and ecosystem report on edge computing. The project provides a glossary with a collection of terms related to edge computing. The project also provides a collaborative market map of how Edge computing technologies and vendors fit together.

2.2.2 StarlingX

StarlingX is a cloud infrastructure software stack for the edge. It uses Docker and Kubernetes and provides a container-based infrastructure. StarlingX puts together a set of well-known open source components and adds to that an openstack configuration tailored for edge-style use cases. On top of that StarlingX adds new software that does the integration of all of the individual pieces including installation, configuration, APIs for management and fault handling.

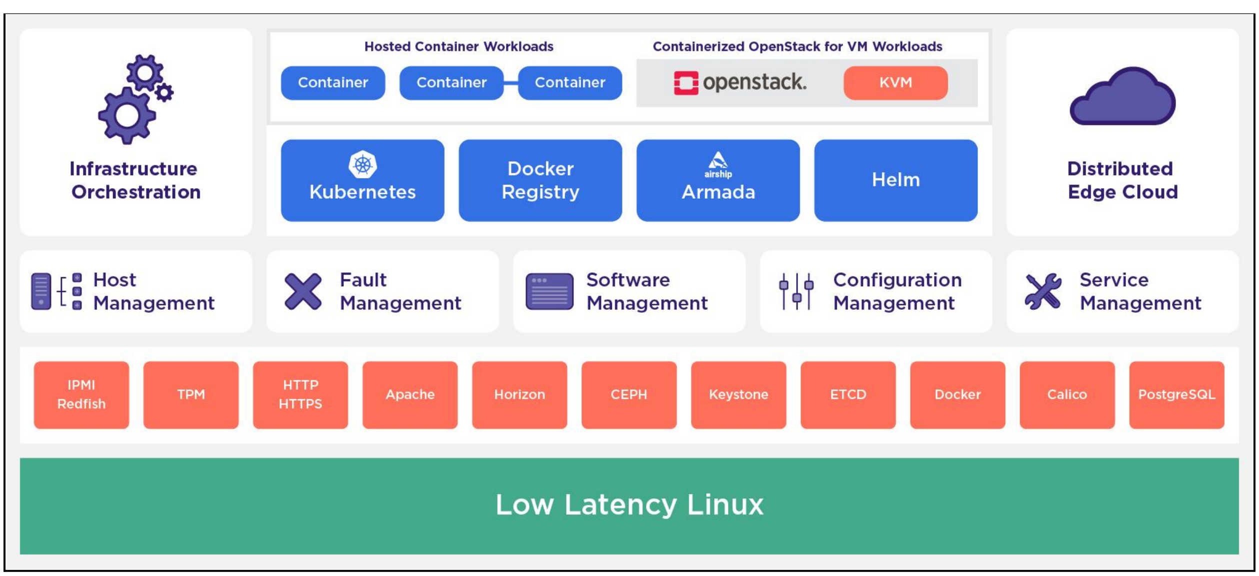

Figure 1. StarlingX architecture (picture from https://www.starlingx.io/).

Starling R4.0 was released in August 2020 and includes support for Kata containers and Time Sensitive Networking.

2.2.3 LightEdge

LightEdge is a standard-compliant implementation of the ETSI MEC Architecture for 4G and 5G networks. The software is released under the APACHE 2.0 licence. Coronado et al [Coronado2020] describes the design and implementation of LightEdge. In the paper, the architecture is evaluated using a latency-sensitive autonomous driving use case in a LTE testbed.

2.2.4 Open19

The Open19 Foundation is an organization established to build a community for a new generation of open data centers and edge solutions. The goal is to create a set of community-driven standards and designs that are customizable, flexible and economical. Much of the focus is on racks and servers and physical equipment. Members include Cisco and Hewlett Packard.

2.2.5 Eclipse Fog05

Eclipse fog05 is an open source project under the Eclipse foundation. The objective of the project is to provide a decentralised infrastructure for provisioning and management of compute, storage, communication and I/O resources in highly heterogeneous systems.

2.2.6 OSF Edge Computing Group

OpenStack is the most widely deployed open source cloud software today. The OpenStack Foundation has an Edge Computing Group. The mission of the group is to identify edge use cases and define the infrastructure systems needed to support distributed applications.
The group has published the white paper on *Edge Computing – Next Steps in Architecture, Design and Testing*.

2.2.7 Containers and orchestration: the Kata, Kubernetes and KubeEdge projects

Containers and lightweight virtualization are key technologies for enabling Edgecloud and distributed applications. The area of container technologies is rapidly evolving. The paper "*The State-of-the-Art in Container Technologies: Application, Orchestration and Security*" [Casalicchio2020], published in January 2020, outlines the current state-of-the-art and research challenges in the area. Examples of Edge Cloud applications are given in Section 4.3. The paper argues that the main challenges for container technologies today are within the areas of performance, orchestration and security. When it comes to performance the biggest challenges are performance prediction, how to best do performance monitoring, and also network I/O throughput improvements. For orchestration the main challenge is run-time adaptation and dynamic control and configuration of multi-container packaged applications on distributed platforms. The main challenges for security are container isolation, confidentiality of containerized data, and network security.

Docker is the de-facto standard for containers today. Bukhary et al [Bukhary2015] evaluate Docker using 4 criteria 1) deployment and termination, 2) resource and service management, 3) fault tolerance and 4) caching. The authors conclude Docker provides fast deployment, small footprint, good performance and good potential for edge computing. The Open Container

Initiative (OCI) was established in 2015 by Docker and other leaders in the container industry. The project creates open industry standards around container formats for runtime execution.

Kata Containers is an open source community with focus on security. The goal with Kata Containers is to combine the benefits of containers and virtual machines. The idea is to provide a container runtime environment based on lightweight virtual machines that has the performance of containers but with additional security through hardware virtualisation.

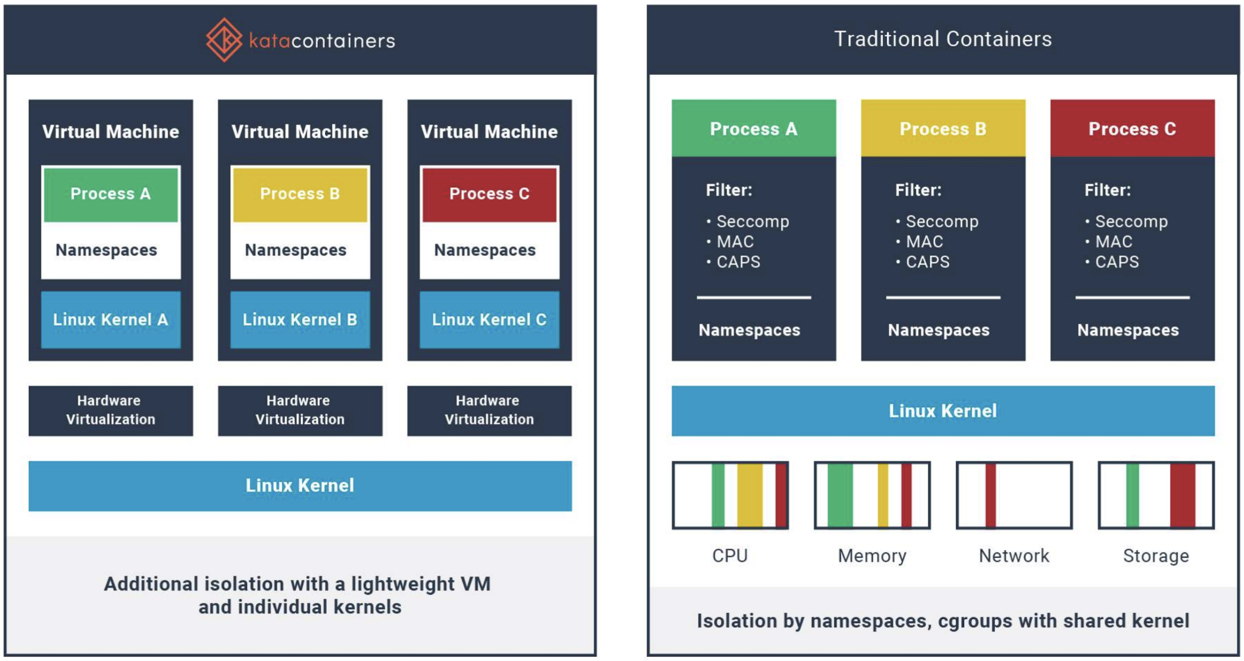

Figure 2. Kata Containers versus traditional containers. Picture from https://katacontainers.io/learn/

The StarlingX cloud software stack supports Kata Containers in it's latest release and Baidu uses Kata Containers in production for Edge computing.

Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications. Kubernetes is today part of many cloud services and Edgecloud projects. Kubernetes was originally developed by Google and then later supported by the Cloud Native Computing Foundation (CNCF). Sometimes known as infrastructure as yaml, Kubernetes is popular. Lightweight Kubernetes, K3s is more appropriate for Edge Computing. Highly available K8s, somewhat confusingly called Micro K8s, distributed by Ubuntu's Canonical is yet another option. Ericsson has their EricKube too. KubeEdge (also supported by CNCF) is an open source system for extending native containerized application orchestration capabilities to hosts at Edge. It is built upon Kubernetes.

An alternative to the existing container approach is to build custom edge-telecom services using a sandboxed secure solution known as WebAssembly. Inspired by browser applications, WebAssembly is a W3C binary format interpreted in an edge enabled VM in this document scope. JIT compilation is in progress, and some rundamentatory networking possible, through WASI and 3rd party initiatives such as lunatics.

2.3 Industry alliances and initiative

2.3.1 The Open Edge Computing Initiative

In the Open Edge Computing Initiative a group of companies are collaborating with Carnegie Mellon University Edge Computing Research Team, with the aim of *"shaping the global ecosystem around edge computing"*. The companies include Microsoft, Intel, Deutsche Telecom, Verizon, Vodafone, MobiledgeX and more. The initiative is running the Living Edge Lab, a real-world edge computing test center for user and technology trials.

"Working Code Trumps All Hype"
-- Prof. Mahadev Satyanarayanan (CMU Edge Computing Research Team)

Mahadev Satyanarayanan (CMU) wrote the paper "The Emergence of Edge Computing" [Satyanarayanan2017] in 2017, explaining the new computing paradigm.

2.3.2 Automotive Edge Computing Consortium

AECC is an industry consortium with focus on network and computing infrastructure for connected car big data. The members include Toyota, Ericsson, Intel, and many more. The consortium is creating use cases and requirements for connected vehicle applications, including high-definition map creation, intelligent driving and remote diagnostics.

2.3.3 Alliance for Internet of Things Innovation (AIOTI)

AIOTI started in 2016 to contribute to the creation of a dynamic European IoT ecosystem and speed up the take up of IoT. Members include large companies, SMEs, research centres and universities. In October 2020 AIOTI published a whitepaper on "IoT and Edge Computing Convergence". AIOTI also organized workshops on IoT and Edge Computing: The Far Edge together with the Horizon 2020 Coordination and Support Action NGIoT, and ARTEMIS (association for actors in Embedded Intelligent Systems within Europe).

References

[Abbas2018] N. Abbas, Y. Zhang, A. Taherkordi and T. Skeie, "Mobile Edge Computing: A Survey," in IEEE Internet of Things Journal, vol. 5, no. 1, pp. 450-465, Feb. 2018, doi: 10.1109/JIOT.2017.2750180.

[Bukhary2015] Ismail, Bukhary & Mostajeran, Ehsan & Karim, Mohd & Tat, Wong & Setapa, Sharipah & Luke, Jing-Yuan & Ong, Hong. (2015). Evaluation of Docker as Edge Computing Platform. 10.1109/ICOS.2015.7377291.

[Casalicchio2020] *The State-of-the-Art in Container Technologies: Application, Orchestration and Security,* Emiliano Casalicchio and Stefano Iannucci, Concurrency and Computation Special Issue paper. *January 2020 https://doi.org/10.1002/cpe.5668*, paper.

[Coronado2020] Coronado et al, "LightEdge: Mapping the Evolution of Multi-Access Edge Computing in Cellular Networks", IEEE Communications Magazine, Volume 58, Issue 4, April 2020.

[DGR/MEC-DEC25TestingFramework] MEC ETSI ISG, “Multi-access Edge Computing (MEC); MEC Testing Framework”. DGR/MEC-DEC25TestingFramework. June 2019.

[Ericsson GFMC-20:000097] “Edge computing and deployment strategies for communication service providers”, White paper, February, 2020.

[ETSI36] “Harmonizing standards for edge computing - A synergized architecture leveraging ETSI ISG MEC and 3GPP specifications”, ETSI White Paper 36, July 2020, ISBN No. 979-10-92620-35-5.

[Hu et al. 2015] Yun Chao Hu et al., *Mobile Edge Computing A key technology towards 5G*. ETSI White Paper No. 11. ETSI. September 2015.

[LightEdge] LightEdge - Transparent Edge Computing Platform, <https://lightedge.io/>, viewed 3 December 2020.

[Mach2017] P. Mach and Z. Becvar, "Mobile Edge Computing: A Survey on Architecture and Computation Offloading," in IEEE Communications Surveys & Tutorials, vol. 19, no. 3, pp. 1628-1656, Third quarter 2017, doi: 10.1109/COMST.2017.2682318.

[Martin 2019] Scott Martin, “What is Edge Computing?” The Nvidia Blog. October 2019.

[Peterson] Larry Peterson et al, “Democratizing the Network Edge”, ACM SIGCOMM Computer Communication Review, Volume 49 Issue 2, April 2019

[Pham2020] Q. Pham et al., "A Survey of Multi-Access Edge Computing in 5G and Beyond: Fundamentals, Technology Integration, and State-of-the-Art," in IEEE Access, vol. 8, pp. 116974-117017, 2020, doi: 10.1109/ACCESS.2020.3001277.

[Reznik 2018] Alex Reznek, “What is Edge?” ETSI MEC blog. May 2018.

[Sabella et al. 2019] Dario Sabella et al., *Developing Software for Multi-Access Edge Computing*, 2nd edition. ETSI White Paper No. 20. ETSI. February 2019.

[Satyanarayanan2017] Mahadev Satyanarayanan, “The Emergence of Edge Computing”, IEEE Computer, Volume: 50, Issue: 1, Jan. 2017.

[Yousefpour2019] Yousefpour et al, All one needs to know about fog computing and related edge computing paradigms: A complete survey, Journal of Systems Architecture, Volume 98, September 2019, Pages 289-330.