POLITECNICO DI TORINO
Repository ISTITUZIONALE

A Disaggregated MEC Architecture Enabling Open Services and Novel Business Models

(Article begins on next page)

19 April 2024

A Disaggregated MEC Architecture Enabling Open Services and Novel Business Models

Gabriele Castellano
Dept. of Computer and Control Eng.
Politecnico di Torino
Torino, Italy
gabriele.castellano@polito.it

Antonio Manzalini
Innovation and Customer Experience
TIM
Torino, Italy
antonio.manzalini@telecomitalia.it

Fulvio Risso
Dept. of Computer and Control Eng.
Politecnico di Torino
Torino, Italy
fulvio.risso@polito.it

Abstract—Network and Service Providers are exploring different exploitation strategies for the Multi-access Edge Computing (MEC), mainly motivated by the opportunities for saving costs and generating new revenues (e.g., through new business models). On the other hand, the overall standardization picture is still very fragmented, delaying or even jeopardizing the real exploitation of MEC; furthermore, current standardization efforts are mainly envisioning a traditional monolithic architecture, with many technological partners but a single administrative domain. This paper argues that a clear separation of IaaS, PaaS and SaaS levels for MEC, together with standardized interfaces, will help accelerating the development of new business roles (e.g., IaaS, PaaS and SaaS providers) and models, possibly replacing the current competition-oriented practices in the telco domain with new forms of cooperation, which are already starting to appear in the IT sector. In this direction, this paper proposes a disaggregated MEC architecture and presents two use cases that show how different categories of resources and services could be provided by infrastructure, platform and software providers in an evolutionary scenario towards 5G.

I. INTRODUCTION

The Digital Transformation triggered by upcoming telecom and ICT technologies will bring a significant impact on current ecosystems, potentially re-designing the roles of Network and Service providers. In fact, the massive softwarization provided by Network Functions Virtualization (NFV), the unprecedented agility guaranteed by Software Defined Networks (SDN), advances in Artificial Intelligence (AI) and the deployment of powerful general-purpose servers at the edge of the (telco) infrastructure, leads to the creation of a *virtual continuum* of resources and services that spans from the edge of the network to remote data centers, encompassing telco nodes, IT servers and potentially even terminals (e.g., smartphones), smart things and IoT devices (e.g., sensors).

In fact, Edge Computing represents an extension of the cloud computing paradigm towards the edge (i.e., between core and access nodes) of a telco infrastructure, aiming at providing better QoS/QoE, optimizing the use of bandwidth, and potentially enabling new services and business models. In this digital transformation, the Multi-Access Edge Computing (MEC) paradigm [1] plays an important role, being it widely accepted as the key technology to meet ultra low-latency requirements as well as to enable a rich computing environment for value-added services closer to end-users.
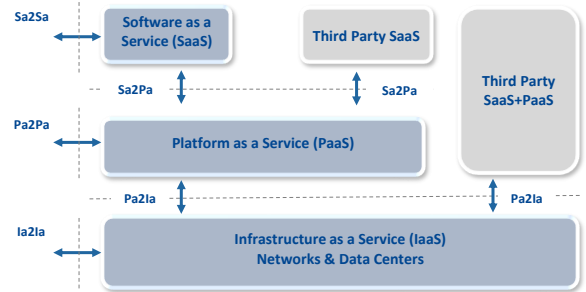

Fig. 1: Layered model for telco 5G actors, inspired by ETSI White Paper "Developing Software for Multi-Access Edge Computing".

Without questioning the importance of MEC in future 5G infrastructures, evidence is rising that current well-established cloud computing models will extend to telco operators not only in terms of individual technologies (e.g., general purpose servers instead of dedicated hardware appliances, deeply reprogrammable networks, agile software micro-services), but also in terms of business models and involved actors [2].

In particular, following the layered model in use in cloud computing [3], [4], we speculate that the business value chain in a typical telco will evolve by originating specialized players such as infrastructure providers (a.k.a., Infrastructure as a Service – IaaS), platform providers (a.k.a., Platform as a Service – PaaS), and software providers (a.k.a., Software as a Service – SaaS). However, novel IaaS, PaaS and SaaS operators may be able to offer resources that (*i*) go beyond the traditional trio of computing/network/storage such as IoT devices, (*ii*) are more heterogeneous and (*iii*) may be present in larger numbers, making a clear evolutionary difference with the existing cloud computing actors. In this respect, each actor will have a clear technical and business role, being in charge of delivering (and selling) richer services to actors that sits higher in the technological stack, which can be seen as a sort of *vertical* interaction toward richer abstractions. In addition, we speculate that, due to the peculiar characteristics of the telco market, additional *horizontal* interactions and novel business opportunities will arise between peers (Figure 1).

In this respect, this paper argues that the current *monolithic* MEC architecture may not be appropriate for future 5G ser-

vices because *(i)* of its unclear splitting of the many required functions between the above mentioned roles (IaaS, PaaS, SaaS), *(ii)* the lack of well-defined software interfaces and *(iii)* the difficulty of enabling new business and service models, and new forms of cooperation between Network and Service Providers and Third Parties. Consequently, without any change to the current MEC building blocks, we propose a novel disaggregated MEC architecture in which clear interfaces between the different actors are foreseen, and show some possible business and optimization opportunities.

II. STATE-OF-THE-ART IN STANDARDIZATION BODIES

Several standardization bodies and fora are addressing MEC and, in general, Edge Computing. Examples include ETSI MEC [5], GSMA [6], TIP [7] — WGs on Edge Computing, OpenFog [8], EdgeX Foundry [9], Open Edge Computing [10], ONF CORD [11], MobiledgeX [12], Akraino [13].

Overall, we are witnessing a fragmentation of efforts in the standardization of MEC/Edge Computing functional and system architectures. On the other hand, all players joining these bodies and fora are recognizing that global interoperability is a must for enabling open services ecosystems and new business models, although most efforts are devoted to north-south interfaces, with little awareness about the necessity of east-west standardization. We argue that the standardization of both east-west and north-south interfaces (e.g., APIs) is crucial to promote innovation and accelerates development of third-party applications, capable of enabling Network and Service Providers to further capitalize on their investments.

III. TOWARD A NEW MODEL FOR MEC

This section defines IaaS, PaaS and SaaS with respect to a MEC scenario, characterizing each actor with a precise role, pursued goals and possible constraints, as well as business and functional interactions with the other actors. A summarized view of these concepts is depicted in Figure 2.

A. Definitions

1) IaaS provider: A IaaS provider virtualizes a set of physical devices, mainly servers and network equipment, offering elementary resources such as CPU, memory, storage, and bandwidth. As shown in Figure 2, resources[1] are provisioned as *virtual execution environments* such as Virtual Machines (VMs) or containers (e.g., Docker) of different sizes. In addition, an IaaS provider has to support *slicing*, i.e., the capability to offer different virtual views of the same physical infrastructure, with strong isolation properties. This represents the key feature to enable *multi-tenancy*, i.e., the capability to support multiple independent users (or *tenants*), possibly associated to different slices. Networking resources include connectivity as well as commonly used functions such as bridges, routers, load balancers, firewalls, tunneling endpoints. In addition, external connectivity (e.g., toward the Internet) must be explicitly advertised in order to enable proper connections to other domains; this represents a key difference with the current cloud-based IaaS model.

2) PaaS providers: PaaS providers offer platform resources and services (also to third parties) for enabling the development of end-users applications and (end-to-end) services. PaaS resources and services include: run-time environment, identity and access management, usage accounting, SLA/QoS management, security, Artificial Intelligence tools, etc.. The resulting software platform requires the creation of an additional programming logic acting as a glue between the different components. A PaaS provider maps high-level deployment requests into low-level IaaS resources, hence consuming infrastructure components (VMs, storage, etc.); this is achieved through a resource orchestrator (Figure 2). In MEC, the design of such a central component needs to take into account *(i)* the heterogeneity of underlying resources, *(ii)* the existence of multiple infrastructure providers and *(iii)* the possible availability of a distributed allocation algorithm. Finally, PaaS should also provide slicing capabilities.

3) SaaS Providers: A SaaS provider delivers turn-key solutions, i.e., ready-to-use applications plus a set of analytics to monitor the state of the service itself; no programming effort is required from its users. Possible examples are video streaming services, content delivery networks, augmented virtual reality, online games. A SaaS provider exploits PaaS facilities to build, deploy and manage the needed topology of *application components* (Figure 2). Such as in IaaS and PaaS, slicing capabilities are required as well.

B. Interaction models

An overall view of all the interactions we identified in MEC is depicted in Figure 2. At a first sight, interactions between IaaS, PaaS and SaaS look similar to the *vertical* interaction model in use in cloud computing: IaaS export resources to PaaS, which deliver end-to-end services used by SaaS providers to create their applications. In case of different interfaces exported by underlying domains, higher-level entities should create their own *adaptation layers*; for instance, a PaaS provider may offer *deperimetrized* services by establishing relationships with multiple IaaS providers.

This paper argues that future MEC players may greatly benefit from an additional *horizontal* interaction model, in which same-level actors can collaborate, enabling each provider to sell also resources owned by its (apparently) competing siblings. This represents a radical departure for most actors in the telco market, which are more used to *competition* than *collaboration*. Although we are aware that this change of paradigm can potentially create disruption in the current ecosystems, we argue that there are both economic and technical reasons why open collaboration may be a better (and forward-looking) option than pure competition even for current telco players. Note that horizontal interactions can be established not only at the connectivity level (e.g. what usually happens between traditional network providers to enable worldwide communication) but, more important, also at the service levels (e.g., between PaaS/SaaS providers).

---

[1]With *resources* we consider both hardware objects (e.g., CPU, network bandwidth, IoT devices), as well as software services (e.g., a database server).
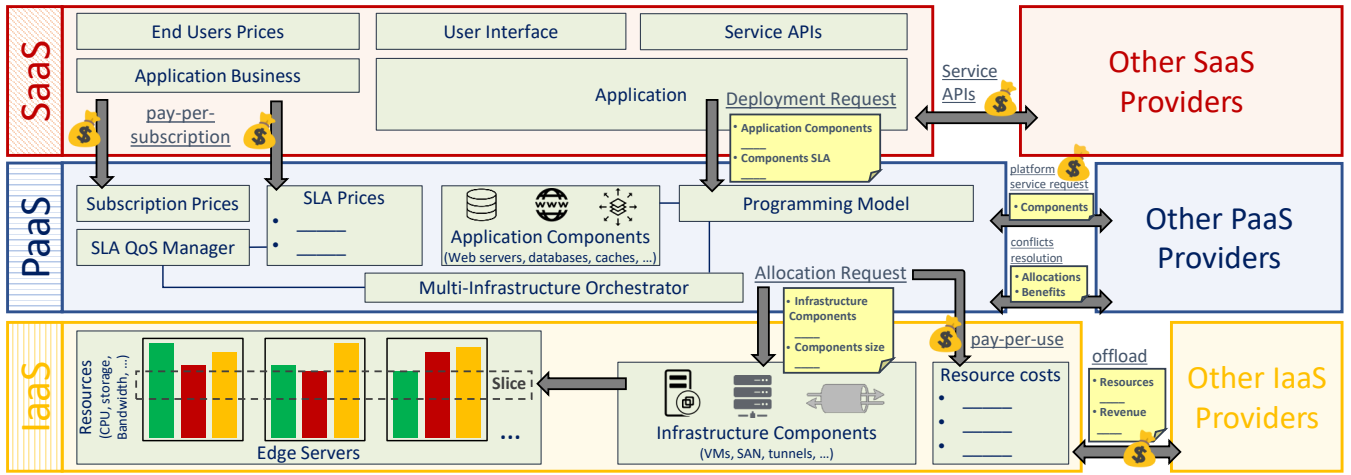
Fig. 2: Overall disaggregated edge architecture with business interactions.

IaaS to IaaS. This interaction follows the well-established model of creating standard interfaces between different operators, which e.g. enables a phone call to cross the boundaries of a single telco. Although future cross-bordering issues may be hidden by creating end-to-end service platforms that establish business relations with multiple IaaS, we foresee several opportunities for an infrastructure provider to collaborate with its peers, enabling to offer a larger infrastructure that may include also (part of) the resources available in partnering IaaS domains. This (*i*) can simplify the operations of a PaaS provider, reducing the complexity required by the interaction with multiple IaaS, and (*ii*) can enable new cooperation strategies between IaaS operators. In fact, they can (*i*) offload services to their IaaS partners in case of resource overload (hence enabling new optimizations) and (*ii*) offer services even in locations in which they have either no access or no economic interest to invest. In particular, this enables an IaaS provider to offer services available on "unconventional" IaaS actors such as enterprise factories or SOHO users, which represent new, tiny IaaS operators. Possible examples of such resources are a set of fog computing nodes in a production plant or an outside temperature sensor / webcam at home. Finally, this may enable further business opportunities for large IaaS providers who can administer the above domain on behalf of their owners, which may not have the technical skills and/or the will to properly operate their infrastructure.

PaaS to PaaS. Future PaaS providers may specialize their offers in specific application domains (e.g., high performance computing, augmented virtual reality), hence possibly requiring their customers (e.g., SaaS provider) to interact with different PaaS actors to create a complex application. In this respect, we foresee the possibility for a PaaS provider to offer services that are in fact provided by another (partner) PaaS. Such as in the previous case, this opens up optimization opportunities for PaaS operators due to the possibility to specialize their offer is specific vertical markets, while buying other (platform) services from their business partners.

SaaS to SaaS. This interaction could be helpful (*i*) to provide a minimum level of compatibility between (competing) applications, and (*ii*) to create bundles of applications. The former would avoid users being forced to use different (but equivalent) applications as multiple instant messaging apps at the same time; the latter enables each SaaS to focus on their core business, hence privileging also in this case a collaborative approach against a competitive one.

C. Interface standardization

For this systemic paradigm change to succeed, standardized interfaces (either de-jure or de-facto) are required. Standards increase the utility of the system by enlarging the number of potential users (Metcalfe's Law) and offer new possibilities for cooperation (horizontal interactions). Given that the northbound interface of a layer can be exploited in either horizontal or vertical interactions, we should define three levels of APIs.

A standard northbound IaaS interface is the initial mandatory step, facilitated by the reasonably clear understanding of current requirements. Existing northbound of open-source IaaS platforms such as OpenStack or Kubernetes can be taken as initial models and possibly extended to accommodate new characteristics such as the presence of heterogeneous computing nodes, diverse network connectivity between nodes (e.g., bandwidth, resiliency), non-negligible network latency due to the geographically distributed infrastructure, and more.

Standardization of higher layers may be more problematic because of the heterogeneity of the services delivered by PaaS and SaaS platforms. However, restricting the scope of the above platforms to the telco domain, we envision the possibility to standardize typical telecommunication services (e.g., Virtual Private LAN Service (VPLS) [14] or Cloud Radio Access Networks (C-RANs) [15]), hence enabling different PaaS/SaaS platforms operating on diverse infrastructure to cooperate for the end-to-end setup of the standardized services.

D. Business models

1) IaaS provider: Infrastructure resources are likely to be provided through a *pay per use* business model. Therefore, an IaaS provider aims at maximizing the amount of resources used by customers over time, e.g., by minimizing the number of physical servers. As shown in Figure 2, a northbound *resource request* will include the set of needed infrastructure components along with the desired resource values (e.g., memory and CPU for each VM, the bandwidth for a network tunnel, etc.). Moreover, the IaaS interface enables to monitor the state of allocated resources (i.e., VM utilization, latency) and to modify allocation at run-time.

2) PaaS provider: A platform provider aims at maximizing the number of customers while minimizing the amount of IaaS resources consumed. At the same time, it has to guarantee any SLA contract stipulated with SaaS providers (e.g., throughput, maximum delay between two application components). The northbound PaaS interface (Figure 2) enables customers to ask for the deployment of application components, along with additional parameters such as logical topology, possible application-specific tunings, SLA constraints. Since platform services are likely to be provided through a *per subscription* pricing model, customers are expected to pay a fixed fee plus the premium cost associated to additional services, such as the desired SLA. Additionally, the interface toward SaaS providers allows also *(i)* state monitoring (e.g., to estimate the current QoE) and *(ii)* run-time components and topology tuning.

3) SaaS provider: As, most likely, also end-consumers will be charged with a subscription, the major goal of a software provider is to maximize the number of users; hence, it has to guarantee a competitive QoE (e.g., throughput, response time). On the other hand, a SaaS needs to minimize the costs of its consumed resources, e.g. requesting the appropriate SLAs and minimizing the number of PaaS providers.

E. Disaggregated MEC

Current standardization for MEC involves the functional definition of the overall architecture and the consequent interfaces between the resulting building blocks. Particularly, the latter aims at guaranteeing the interoperability between different implementations and, potentially, different vendors, which is only one of the requirements in case of a scenario in which Infrastructure, Platform and Software services may be provided by different actors. In fact, in this case the interface must support additional parameters such as authentication, accounting, billing; the current monolithic architecture, instead, assumes that all the functional blocks are under control of the same organization.

We propose the disaggregated MEC architecture depicted in Figure 3, which shows a possible splitting of functionalities between different IaaS, PaaS and SaaS providers. In addition, we highlight some interfaces (Mp2, Mm6) that must be extended to support cross-actor interactions, enabling the MEC platform to run on an infrastructure owned by a different entity, and even on different infrastructure domains. Furthermore,
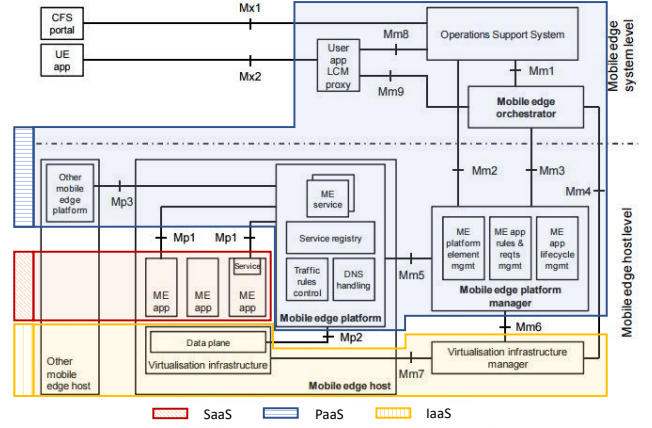

Fig. 3: Disaggregated MEC architecture.

interface Mp1 enables a third-party software provider to install its applications on top of the MEC platform.

The standardization of the above interfaces will have a profound impact on both technological and business sides. With respect to the former, the different pieces of the architecture can be sold (if we refer to technological providers) and/or operated (with regard to network providers) by different actors, which can be different business units of the same company or different companies. On the business side, this would open up the market to multiple specialized actors, which can establish either competition or cooperation relationships. Finally, we expect this to bring a new breed of novelty to customers, which can experience new services and innovative offers thanks to breakage of current monolithic network providers, replaced by more specialized and possibly competing actors.

This would transform the relatively slow world of network and service providers into a fast growing and innovative area, similar to personal computers and smartphones. In fact, the splitting of concern between different (vertical) actors is considered one of the keys of the extraordinary evolution of the above markets over time, which was possible by unleashing the power of independent developers that were continuously enriching the original platform. Similarly, this vision aims at an innovative evolution of the concepts of Service Delivery Platforms and Service Layers of former times. Specifically, IaaS-PaaS-SaaS decoupling together with standardization (either de-jure or de-facto) of the related interfaces will overcome static and vertical service silos that delay open innovation. That "openness" will enable multiple interactions of providers at different levels, will eases the role of apps developers, allow "deperimetrization" of services and create new business opportunities for all players of the new ecosystems.

IV. USE CASES

As an example, we now present how two novel possible use cases can be mapped to the disaggregated MEC model, identifying actors and analyzing their interactions both from a business and technical point of view.

A. Everywhere in the city

In this first use-case we provide to each person the set of digital services more appropriate in each given location. For instance, in a supermarket we provide information about available offers; at work we offer fast access to shared printers and servers; at home present a dashboard to control ambient-assisted living and smart appliances, heating, and more. Other examples are the provisioning of infotainment services to both car drivers/passengers and to pedestrians in smart roads and city hot spots (e.g., commercial areas, stadium, stations, etc.).

We speculate there might be one or more telco IaaS providers covering the overall city, while the stadium sets-up dedicated IT equipment to be used in case of live events, and the supermarket shares part of its IT infrastructure for serving local people. All the above actors may collaborate to provide an ubiquitous service through a logically partitioned infrastructure. A disaggregated MEC can enable end-to-end services, being them either provided through different infrastructures or platforms. The pervasive distribution of MEC IaaS and PaaS will allow to improve the users experience thanks to latency reductions and zero service interruption on mobility. From a business perspective, a telco may wish to play the role of IaaS provider and/or PaaS provider; in the former case, it will need to engage business relationships with third parties playing the roles of MEC PaaS providers (e.g., municipality, supermarket, stadium...) owning MEC platforms.

B. Enterprise connectivity services

This second use case envisions an enterprise willing to offer more integrated connectivity services to its employees. The most common current options are to buy/install services such as VPN, and/or buy a "private mobile network" to enable its users to communicate at convenient prices, possibly with internal (short) phone numbers. In case of telco services provided by distinct IaaS-PaaS-SaaS actors, the above enterprise may have several additional options. In fact, it can easily become a "virtual operator", providing its remote users with native network services, being them xDSL connections or direct data mobile connectivity (through enterprise-branded SIM cards), with the assurance that its remote users will be always securely connected to the corporate network without having to install/launch any VPN software. The "virtual operator" service can be either created by buying the elementary components from a PaaS provider and adding the additional logic to create the requested service, or by simply buying a turn-key software from a SaaS provider. The PaaS provider, from its side, will be in charge of establishing the proper business agreements with multiple IaaS providers (e.g., the IT infrastructure in the main corporate site, which represents the main infrastructure to connect to) to enable mobile users to take advantages of the above services whatever physical infrastructure they are in.

V. CONCLUSIONS

MEC is expected to play a key strategic role in the Digital Transformation towards 5G. As a matter of fact, network and service providers are exploring different strategies to introduce and exploit MEC concepts, mainly motivated by the potential opportunities for saving costs and generating new revenues.

Although several standardization bodies and fora are targeting MEC and, in general, Edge Computing, current solutions are still lacking end-to-end interoperability, which is must for enabling open ecosystems. This paper argues that a clear architectural *decoupling* of IaaS, PaaS and SaaS for MEC, and then for 5G, represents an evolutionary step of the digital transformation capable of enabling new roles and business models. Furthermore, this may blur the borders between current OTT operators, which are mostly offering datacenter-only services, and network providers, which own the network infrastructure that connects datacenters to the customers. In this respect, network providers represent the most suitable actor to offer also edge-based resources such as IoT and domestic/enterprise-owned IT infrastructures at the edge of the network, hence positioning themselves in a stronger position when competing with current OTT actors.

This paper addresses the different resources and services for IaaS, PaaS and SaaS providers and provides examples in terms of optimization and monetization. Eventually the paper proposes to pursue joint efforts of network and service providers and technology vendors (of MEC hardware and software solution) in industry bodies and standard fora for the MEC IaaS and PaaS architecture and interfaces standardization. In fact, this model is likely to boost open services ecosystems with new business models and forms of cooperation between network and service providers with third parties.

REFERENCES

[1] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella, "On multi-access edge computing: A survey of the emerging 5g network edge cloud architecture and orchestration," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 3, pp. 1657–1681, 2017.

[2] B. Sayadi and L. Roullet, "5g: Platform and not protocol," 01 2018. [Online]. Available: https://sdn.ieee.org/newsletter/january-2018/5g-platform-and-not-protocol

[3] B. P. Rimal, E. Choi, and I. Lumb, "A taxonomy and survey of cloud computing systems," in *INC, IMS and IDC, 2009. NCM'09. Fifth International Joint Conference on.* Ieee, 2009, pp. 44–51.

[4] M. Litoiu, M. Woodside, J. Wong, J. Ng, and G. Iszlai, "A business driven cloud optimization architecture," in *Proceedings of the 2010 ACM Symposium on Applied Computing.* ACM, 2010, pp. 380–385.

[5] S. Kekki *et al.*, "MEC in 5G networks," Tech. Rep., 06 2018.

[6] GSMA. [Online]. Available: http://gsma.com

[7] (2016) Telecom infra project. [Online]. Available: http://telecominfraproject.com

[8] Openfog consortium. [Online]. Available: http://www.openfogconsortium.org

[9] EdgeX Foundry Architectural Tenets. [Online]. Available: https://docs.edgexfoundry.org/Ch-Intro.html

[10] Open edge ccomputing. [Online]. Available: http://openedgecomputing.org

[11] A. Al-Shabibi and L. Peterson, "Cord: Central office re-architected as a datacenter," *OpenStack Summit*, pp. 1–38, 2015.

[12] Mobileedgex. [Online]. Available: mobiledgex.com

[13] Akraino Edge Stack. [Online]. Available: https://www.akraino.org/

[14] K. Kompella and Y. Rekhter, "Virtual private lan service (vpls) using bgp for auto-discovery and signaling," Tech. Rep., 2007.

[15] C. M. R. Institute, "C-ran: The road towards green ran," 10 2011.