

Fog Computing and Mobile Edge Cloud Gain Momentum

Open Fog Consortium, ETSI MEC and Cloudlets

Version 1.0

Guenter I. Klas

Nov 22, 2015

Abstract

Given the announcement of the Open Fog Consortium on 19 Nov 2015, this paper provides an introduction to Fog Computing and puts this concept into a wider context. Fog Computing is compared to similar and related approaches like Mobile Edge Computing and Cloudlets. We illustrate what all three approaches are good for. After observing that the mentioned three approaches to distributed cloud computing for Internet of Things and Tactile Internet have emerged relatively independent of each other, we stress that there is room for synergy. Finally we point out a number of areas that require further work and research in order to get concepts like Fog, Mobile Edge Computing and Cloudlets adopted by the industry as extensions to the classical Cloud computing.

Tags: Fog computing, fog node, 5G, cloud computing, cloudlet, edge computing, mobile edge computing, edge cloud computing, open edge computing, intelligent cloud, OEC, MEC, ETSI MEC, low latency, predictable latency, robotics, tactile internet, Internet of Things, IoT, Victor Bahl, Cisco, Microsoft, OpenFog, Open Fog, OpenStack, distributed cloud applications, real-time analytics.

1. What is Fog Computing?

Cisco has been a promoter of fog computing for a while. As engineers from Cisco say in [1], fog computing extends the cloud computing paradigm to the edge of networks, in particular wireless networks for the Internet of Things (IoT). If big datacentres and high-capacity network backbones correspond to our city centres with skyscrapers, the rural suburbs correspond to the “edge” of networks. Seen from the centre, the edge is “further out there”, in the twisted maze of access networks, wireless macro and small cells, in general “where the devices live” ... and where the billions of things of the IoT are consuming control commands and sending data – towards the Cloud.

The name *Fog Computing* comes from the following analogy: When you get up early in an autumn morning, you see the first sunrays shining through fog. Above the fog, there may well be a few clouds sprinkled in the blue sky. The fog is closer to you than the clouds. Regarding distance, an Internet of Things device is closer to a fog computing platform than to most of the typical large-scale datacentres (those we find today are arranged as one or two per ‘region’ across the globe, see e.g. Microsoft Azure regions).

Characteristics of fog computing are as follows [1]:

- a) Fog computing nodes are typically located away from the main cloud datacentres, at the edge.
- b) Fog computing nodes are *wide-spread* and geographically available *in large numbers*.
- c) Fog application code runs on fog nodes as part of a distributed cloud application.
- d) Cloud computing on fog nodes enables *low* and in addition *predictable latency*.
- e) Fog computing nodes provide applications with awareness of device geographical location and device context.
- f) Fog computing nodes can cope with mobility of devices, i.e. if a device moves so far away from a fog node that the service latency becomes too poor, the fog node can redirect the application on the mobile device to associate with a new application instance on a fog node that is now closer to the device.

- g) Fog nodes offer special services that may only be required in the IoT context (e.g. translation between IP to non-IP transport).
- h) Fog nodes are typically accessed by devices over wireless networks.

2. Which application areas are typically in the scope of fog computing?

The areas mentioned most often are real-time applications, connected vehicle (V2V, V2X), smart grid, smart cities, wireless sensor networks combined with actuators [1], pipeline monitoring, connected rail, smart traffic light systems, wind farms, closed loop control of industrial systems, orchestration of a number of autonomous yet coordinated components (e.g. coordination of many distributed closed loop control systems at a wind farm), and applications in the oil and gas sector [2].

3. What are the key arguments Cisco uses to underpin the need for fog computing?

A key argument goes as follows: The deployment of IoT applications in a 2-tiered way doesn't meet the requirements related to low latency, mobility of the "things" and location awareness [1]. 2-tiered deployment means that the first part of the IoT application runs on the "thing", i.e. the IoT device, and the second part runs in a hyperscale datacentre (e.g. Microsoft Azure).

Thus, the solution to this problem is a multi-tiered architecture (with at least 3 tiers) whereby an IoT application is deployed as follows: a part on the "thing" (e.g. a car), a 2nd part on the fog platform (e.g. a roadside cabinet or a router in a wireless access network or an LTE base station), and in the case of three tiers a final 3rd part in a datacentre of the main cloud (e.g. Amazon EC2).

4. Is the fog computing concept from Cisco the only game in town?

The answer is no. Similar ideas have been developed by various other parties and the list of those might be rather long. Suffice to state three other efforts for the purpose here:

- a) The European Telecommunications Standards Institute ETSI launched an industry specification for **Mobile Edge Computing** (MEC) in Sept 2014. This group is developing a system architecture and standardising a number of APIs for essentially mobile edge cloud computing. See [5] for the main site and [9] for their wiki.
- b) A team at Carnegie Mellon University (CMU) has developed **Cloudlets**. As stated in [11] a cloudlet represents the middle tier of a 3-tier hierarchy: 'mobile device – cloudlet - cloud. A cloudlet can be viewed as a "data center in a box" whose goal is to "bring the cloud closer"'. CMU have pioneered particular research in this space and implemented various mechanisms as open source code which is e.g. available at [12]. Recent publications that underpin the usefulness of cloudlets are [13] and [14]. An in-depth technical introduction is available in [16].
- c) Microsoft Research, with Victor Bahl, has socialised the concept of **micro datacentre** as an extension of today's hyperscale cloud datacentres (as Microsoft Azure) in order to meet new application demands like lower latency and new demands related to devices (e.g. lower battery consumption). For an interview with Bahl, see [17] and for a presentation about micro datacentres see [18].

The immediate conclusion is: fog computing is not the only game in town, not the only flower in the flower bed. And it's worth to look closer at some of the differences.

5. How does Fog Computing differ from ETSI MEC and Cloudlets?

For an introductory paper to mobile edge cloud computing and further important references, see e.g. [19]. That paper addresses both ETSI MEC and Cloudlets.

A **highlight of ETSI MEC** is the objective to deliver a standardised mobile edge computing architecture and industry-standardised APIs for 3rd party applications.

A **highlight of Cloudlets** is that much effort has been put into i) defining a system and creating algorithms that support low latency edge cloud computing and ii) implementing the relevant features in open source code as extensions to OpenStack cloud management software [12]. Two key features are near-real-time, just-in-time provisioning of applications to edge nodes and handoff of virtual machine images from one edge node to the next edge node once a device has moved away from its first edge node and has come closer to a new edge node.

A **highlight of Fog Computing** is that it assumes a completely distributed, multi-layer cloud computing architecture, where there are billions of devices (as part of the IoT), lots of Fog and finally the main hyperscale cloud datacentre. A single, particular Fog application is typically distributed across such devices, across multiple fog components inside the Fog and some datacentre cloud. Thus, the Fog itself is made up of fog cloud computing components embedded in nodes in different network tiers, e.g. in the radio access network, the multi-service edge, the core network (on IP/MPLS routers/switches, at mobile packet core network gateways etc.). See e.g. Fig 4 in [2].

The authors from Cisco in [2] also stress the need to define the virtual network topology for each (fog) application. In comparison, neither the papers available on Cloudlets nor ETSI Mobile Edge Computing highlight the important role of software defined networking to implement and provision such virtual network topologies that connect devices with multiple fog components and a cloud datacentre (see [2] p12).

Before we look into more detail of Fog Computing, Mobile Edge Computing and Cloudlets, let's contrast the roles of the different hierarchical layers, from devices to the hyperscale datacentres.

6. What are the roles of device, edge (fog) layer and main cloud (datacentre)?

In Table 1, we list a set of examples for software functions that are typically allocated to a device or an edge computing platform or a main cloud.

Device	Edge or fog layer	Main cloud (datacentre)
Provides user interface (I/O, rendering of output)	Hosts as 3 rd party apps "network" functions like video acceleration	Provides data storage (long permanence)
Hosts micro-control of actuators	Hosts middleware like registry for edge applications, inter-edge-application communications, services that provide access radio network information. These may be considered components of PaaS at the edge.	Provides human to machine interface for overall application management (e.g. dashboard, deployment, provisioning)
Hosts micro-control of on-board sensors	Offers APIs to 3 rd party applications executed at the edge	Hosts visualisation and reporting for operations
Hosts local compute, storage, network stack	Hosts compute-heavy parts of the overall application (e.g. object recognition, motion classification)	Provides off-line, batch data analytics software (maybe real-time analytics as well)
Others	Hosts real-time analytics software	Hosts machine-learning software
	Hosts latency-sensitive control-loop software	Serves queries from device with response time > 100ms ¹⁾
	Issues control commands to devices and actuators	Hosts enterprise integration components
	Collects M2M/IoT data incl. sensor data	Others

	Filters data (to consume locally or to send to the main cloud)	
	Serves queries from device with required response times in a range [1ms to say 100ms] ^{1) 2)}	

Table 1: Roles of device, edge cloud / fog and main cloud datacentre

Notes:

- 1) Response time is measured on application level from first request to first response (i.e. round-trip). The cut-off point (upper end of the latency range for the edge) is a choice made here. Various experts will have different views on this. Closed loop systems incl. high-speed remote robot control require lower response times (e.g. <10ms), apps working with cognitive assistance might work with a worst-case response time of e.g. 30ms.
- 2) The paper in [2] about Fog Computing states a maximum reaction time permissible for a Smart Traffic Light System of 10 ms.

7. A comparison of Fog Computing, ETSI MEC and Cloudlet

In the following, we provide a comparison of the three industry approaches in Table 2.

Item	Fog Computing [1], [2]	ETSI MEC	Cloudlets
Uses a virtualised IaaS platform?	Yes	Yes	Yes
Allows multi-tenancy of apps at the edge?	Yes	Yes	Yes
Is located between end device and main datacentre?	Yes	Yes	Yes, however a cloudlet can run also directly in a 'device' (e.g. a car, or aeroplane)
Is an extension of the Cloud (complements it)?	Yes	Can be, but doesn't need to be	Yes, typically is
Its computing platforms are geographically distributed?	Yes	Yes	Yes
Can be physically co-located with e.g.:	Access points, base stations, traffic aggregation points, routers, switches, gateways	Same	Same
Is inspired by the needs of IoT services that require distributed computing and storage?	Yes	Less so. Greater focus on enabling apps that require low latency and device context (e.g. radio access network)	Less so, however takes inspiration from Tactile Internet
Assumes it is mostly used together with wireless access?	Yes	Yes	Not necessarily (can equally be used with fixed access)

Item	Fog Computing [1], [2]	ETSI MEC	Cloudlets
Enables low latency?	Yes	Yes	Yes
Enables low jitter?	Yes	Yes	Yes
Enables applications at the edge to be context-aware?	Yes	Yes (in particular aware of the radio access context of a device, e.g. available bandwidth in a cell)	Not intrinsic to Cloudlets, but can be added
Offers support for geographical mobility of application on device?	Yes	Yes	Yes
Requires platform services to be federated across domains of different edge node ownership and providers?	Yes (thus calls for interoperability, therefore the new organisation Open Fog Consortium)	Yes (thus done in ETSI as standards-defining organisation)	Yes (thus the desire to incorporate key APIs into the open source project OpenStack)
Has particular focus on on-line analytics and interplay with the Cloud?	Yes	No	Not specifically
Supports applications typically deployed in an N-tier hierarchy?	N = 3 or more (distributed Fog infrastructure [2])	N = 2 or 3 (2: device + one edge location) or (3: device + one edge location + main cloud)	N = 3 typically
Enables deployment in different versions, including ruggedized for outdoor use?	Yes (e.g. a fog node at a traffic light)	Yes (e.g. together with an LTE base station)	Yes
Is near-real-time interaction amongst same apps on different edge nodes a key consideration?	Yes, in the sense that inter-fog node communication is supporting a fully distributed application (e.g. communication between smart traffic lights). See [2] Sect 3.1.2 for more explanation.	Partly (so far only to support mobility of a device when the device disassociates from edge node 1 and associates with a new edge node 2).	Partly (similar to left column), to support “handoff” of a virtual machine image from edge node 1 to edge node 2 to get the edge app closer to a moving device.
Stresses the need for efficient communication between edge nodes?	Yes	No	No
Provides APIs for provisioning and monitoring virtual resources for compute, storage, network?	Yes: The “Fog Abstraction Layer” provides such APIs [2]. Foglet software agents use such APIs and constitute a distributed fog orchestration framework. [2], Sect. 6.3.1.	Yes, currently assumed this is done via a Mobile Orchestrator (borrowing from ETSI NFV MANO and service orchestration for network function virtualization)	Yes, such APIs are exposed via OpenStack and extensions to OpenStack as proposed at the OpenStack Summit in Tokyo, Oct 27-30, 2015 [15]

Item	Fog Computing [1], [2]	ETSI MEC	Cloudlets
Provides life-cycle management of distributed cloud apps?	Yes, via the “Fog Service Orchestration Layer” [2]	Yes, currently via Mobile Orchestrator and OSS/BSS of the telecoms network operator	Partly specified
Aims at supporting different use cases from multiple vertical industries?	Yes (e.g. smart cities with smart traffic lights, energy (wind farms))	Yes (e.g. security industry, content delivery industry)	Yes (e.g. health sector, security sector, consumer discretionary with cognitive assistance)
Has been originally promoted by?	Cisco	A group of 6 companies who founded ETSI ISG MEC: Nokia, Huawei, IBM, Intel, NTT DoCoMo, Vodafone	Prof. Satya, Carnegie Mellon University [21], later supported by various companies including Intel, Huawei, Vodafone
Which drivers inspired the concept?	IoT, Wireless Sensor and Actuator Networks (WSAN)	NOKIA Liquid Apps, Telecom use cases like content acceleration, augmented reality etc. MEC as stepping stone towards 5G mobile networks.	Research into distributed cloud and requirements stemming from Cognitive Assistance applications, see e.g. Gabriel architecture

Table 2: Comparison of three approaches: Fog, Mobile Edge Computing, Cloudlet

Next, a few statements which are more related to the organisations that drive above three approaches and a comparison of some of their more important goals are provided in Table 3.

	Open Fog	ETSI MEC	OEC (Cloudlets)
Purpose	To drive industry and academic leadership in fog computing architecture, testbed development, and a variety of interoperability and composability deliverables that seamlessly leverage cloud and edge architectures to enable end-to-end IoT scenarios. [4]	<ul style="list-style-type: none"> The work of the MEC initiative aims to unite the telco and IT-cloud worlds, providing IT and cloud-computing capabilities within the RAN (Radio Access Network). The ISG MEC will specify the elements that are required to enable applications to be hosted in a multi-vendor mobile-edge computing environment. [5] 	<ul style="list-style-type: none"> Leveraging Cloudlets as enabling technology. Driving the necessary technology (e.g. extensions to OpenStack, KVM, QEMU). Driving prototyping of applications that leverage edge cloud computing, pushing the boundaries and demonstrating benefits. Driving the eco system development for Open Edge Computing. Engaging with target service industries/sectors through demonstrators and joint projects. Engaging with the relevant developer communities, seeking feedback and driving acceptance of edge cloud computing. Synchronising work with other efforts incl. ETSI ISG MEC and OPNFV [34]

	Open Fog	ETSI MEC	OEC (Cloudlets)
Unique highlight	<ul style="list-style-type: none"> Plans to further develop an architecture framework for Fog, to influence standardisation and to promote open fog computing implementations Drives this from an IoT perspective and the need for distributed cloud computing for IoT. 	<ul style="list-style-type: none"> Creates architecture framework and standards for APIs to 3rd party applications located at the edge. Addresses how edge applications can be enriched through real-time context information from the radio access network. 	<ul style="list-style-type: none"> Offers open source code for Cloudlets (extension to OpenStack) as ecosystem enabler [12]. Targets any industry that benefits from low-latency edge cloud computing, whether products fall into IoT, Tactile Internet (including systems with haptic feedback), 5G mobile networks, web content delivery, or on-line gaming etc. Scope wise limits itself to focus on a few key enablers, thereby not constraining the spectrum of use.
Goals			
Create an architecture framework	Yes, for Fog	Yes, for MEC	Less required, as Cloudlets are an enabler that can be used for Fog computing and for Mobile Edge Computing.
Have open forum to share ideas	Yes, for open fog computing	Less applicable. To participate, must either be ETSI member or sign ETSI MEC Participation Agreement to respect the ETSI IPR policy	Yes, for Cloudlets. http://forum.openedgecomputing.org/
Create educational material	Yes, for Fog	Yes, see ETSI MEC Industry Enablement Group [6]	Yes, for Cloudlets
Promote implementations	Yes, of Fog Computing	Yes, of ETSI-compliant Mobile Edge Computing	Yes, of Cloudlets
Influence standards	Yes	Yes, is part of a standards defining organisation itself	Yes (e.g. influencing ETSI MEC)
Influence relevant open source communities	Not stated in activities list at http://www.openfogconsortium.org/about-us/ as per 22 Nov 2015	No	Yes, in particular OpenStack
Run a technology testbed	Yes, at Princeton University	No, however members are encouraged to run tests by creating Proofs of Concept [7], [8]	Yes, at Carnegie Mellon University
Host plugfests	Yes	Supported by the ETSI Center for Testing and	Not announced yet.

		Interoperability (CTI). See [7].	
Reach out to vertical industries	Assumed	Yes. See ETSI MEC Industry Enablement Group [6]	Yes, see purpose.
URL	http://www.openfogconsortium.org/	http://www.etsi.org/technologies-clusters/technologies/mobile-edge-computing	http://openedgecomputing.org/about-oec.html

Table 3: Aspects related to three groups advocating Fog, Mobile Edge Cloud and Cloudlets.

Though there are differences in focus for multiple reasons, all three efforts share a similar vision. This vision is driven by an anticipated future which is characterised by the Internet of Things, Internet of Everything, Tactile Internet, and existence of appropriate wireless connectivity, whether low data rate plus minimal power consumption at a device (e.g. NB-IoT) or super-high bandwidth for ultra-low latency (e.g. 5G).

This perspective inevitably leads to the conclusion that the current paradigm of bigger and more consolidated, more hyperscale cloud computing datacentres will fall short of future industry needs. What's required is indeed an extension of cloud computing to the edge of networks.

Various players have recognised this early on, as reflected in the existence of concepts like Fog Computing, Mobile Edge Computing and Cloudlets (Open Edge Computing). The initiation has come less from the sides of the big public cloud providers like HP, IBM, Google, Microsoft, or Amazon.

However, standing out and therefore worth to note is the vision as shared by Microsoft Research (Victor Bahl) about expanding hyperscale datacenters with a number of much smaller satellites, called micro datacenters (or mcDC). See [18] for the conference paper and [17] for an interview.

8. What can Fog Computing, ETSI MEC and Cloudlets all enable?

A number of use case areas were already mentioned above when we introduced Fog Computing. Implementations following all these three approaches can play a role e.g. in:

- Connected vehicles, V2V, V2X. Automotive safety services (like ice on motorway real-time warning, platooning, coordinated lane change manoeuvres etc.).
- Services in infotainment (e.g. in automotive).
- Smart city system components like smart traffic lights which go beyond what's available in some countries today. For example the traffic lights are able to send warning signals to approaching (autonomous) vehicles [1]. The paper [2] explains: Smart traffic lights “detect the presence of pedestrians and cyclists crossing the street. The STL also issues “slow down” warnings to vehicles at risk to crossing in red.”
- Big data and analytics for sectors like industrial: real-time analytics at the edge, long-term analytics in the main cloud, for purposes like predictive maintenance and others.
- Smart grid, closed loop system (see also next).
- Closed loop systems where a degree of central processing occurs at the edge. For such systems, system stability and avoidance of oscillatory behaviour is important. In this case, the motivation for fog or edge cloud computing is not only low latency for rapid responses, but low and predictably low jitter.
- Efficient operation of wind farms: semi-autonomous controllers at each wind turbine are orchestrated on the level of the wind farm by software at the edge [2].
- Robotics for various sectors including assisted living, remote diagnostics etc.
- Any sector that makes use of wireless sensor systems, whether in oil & gas industry or building industry.

- And lots of others as documented in various literature.

Note: Illustrative, detailed descriptions of a Smart Traffic Light System (STLS) use case and a wind farm use case are provided in [\[2\]](#).

9. What about the aspirations regarding application deployment at the edge?

I suggest distinguishing between two edge cloud/fog environments: curated environment versus non-curated environment. The aspirations of both are different and the environments take into account the expectations of application developers for edge-apps or foglets.

Curated environment: It requires a 3rd party edge-application developer to closely work together with the infrastructure vendor who delivers the edge cloud/fog platform and the party who operates this platform. The edge application will be vetted for security, extensively tested and possibly integrated in a customised way (apart from it possibly making use of standardised APIs as e.g. from ETSI MEC). An example scenario: A developer creates video acceleration software which is deployed on the edge cloud/fog computing platform of a router manufacturer, whereby the router is deployed in the access network of a mobile (cellular) network operator. The expectations of the 3rd party app developer are met, as that party is used to integration work with telecoms companies.

Non-curated environment: It enables deployment of edge applications as easy as it allows deployment of applications into any of the main clouds today (Amazon EC2, Microsoft Azure, IBM, HP, Google, ...): Fast, easy, efficient. This would be the way how an online games company or an augmented reality software company would prefer deploying their applications e.g. in order to benefit from reduced latency and the existence of fog. The challenge here is to create a seamless deployment process for edge-enabled cloud applications whereby the process stretches across domain boundaries (e.g. an application gets deployed across Amazon EC2 plus 100 edge cloud nodes from a different provider).

Further attention needs to be paid to create relevant application deployment and resource provisioning processes and tools, which leverage best practice from the IT side (Cloud) and from the Telecom/ISP Network side (Network Function Virtualization).

10. Similar ideas addressed from different angles: fog, mobile edge cloud, cloudlets

It's worth noting that similar concepts have been pushed by different players with different backgrounds and of course different business interests. Vendors who are key players in the Internet router market would argue that their routers and switches are the perfect platforms to host edge cloud/fog node capabilities (e.g. Cisco, Juniper). Telecoms equipment vendors in turn may argue their equipment (e.g. 4G, 5G base station, gateways etc.) are the natural physical locations of hosting edge computing capabilities.

Different parties bring their own legacy and industry knowhow to bear to solve a very similar challenge. E.g. when it comes to provisioning of virtual and physical resources for compute, storage and networking, the Cisco paper on Fog computing [\[2\]](#) introduces its own terminology (foglet software agents, distributed databases, policy managers, capability engines, ...), whilst multiple players amongst the group of companies pushing ahead with ETSI MEC urge to maximise synergy with mechanisms anyway being developed for network function virtualization (NFV), including the orchestration mechanisms that have been defined by ETSI ISG NFV [\[10\]](#).

From the comparison provided above, one may argue there is a degree of synergy available between the three concepts. It appears important to recognise the truly innovative ideas that have emerged as part of the different approaches, whether Fog Computing, Mobile Edge Computing or Cloudlets and to create solutions that leverage those innovative ideas.

11. Conclusions and what we need to do more of

- a) Bringing into the discussion of Fog and Edge Cloud Computing some of the big players in public and hybrid clouds (and through that indirectly their customers in various industry sectors). Why?

The places where cloud computing product roadmaps leading into the future get discussed are the bilateral meetings between clients and vendors (e.g. between Rockwell Automation and Microsoft Azure [22], [23]). This is the realm of business development and it represents a safe environment (where if need be an NDA can be put in place). Companies like Rockwell Automation represent the end users in vertical industries. It's unlikely that such companies directly engage with groups like ETSI ISG MEC or OEC. Though ETSI ISG MEC has an Industry Enablement Group that develops collateral to reach out to vertical industry sectors, ETSI MEC cannot bring a Cloud computing product roadmap to the table which would be of short- to mid-term interest to the client. The same challenge exists for OEC and possibly for Open Fog.

How to bridge the gap between "the concept" (like Fog, Mobile Edge Computing, Cloudlet) and the industrial corporate client (like Rockwell Automation)? One way forward is to sell the concept to industry leaders in Cloud computing, and use those leaders' established business development channels, business relationships and credibility in Cloud computing to introduce the new technology to the ultimate clients. The news that Microsoft has co-founded Open Fog is therefore highly important. Another way is for players who own the assets in the emerging edge cloud/Fog to team up with tier 1 cloud computing providers and vendors to jointly address the market opportunities in vertical sectors.

Fog Computing, ETSI MEC and Cloudlets all have to gently insert themselves (directly or indirectly) into the discussions that already forge the future of Cloud computing for IoT. Take the case study of Thyssen Krupp and Microsoft Azure [24]. Cloud computing at the edge (in a Cloudlet, in a Fog node, on a MEC server) may provide benefits on top of what the public cloud providers have already on their roadmap when they transform the Cloud into the "Intelligent Cloud" [31].

The steps already being undertaken to transform the hyperscale public clouds into Intelligent Clouds ready for the IoT include:

- Expanding the footprint of the public cloud datacentres into more 'regions' (e.g. the number of Microsoft Azure regions is growing [25], [26], however the number is in the low double digits. Compare this to the vision of having large numbers of fog nodes widespread across the world as in Fog Computing).
- Adding cloud gateways that are located close to the IoT data sources and help inject data into the main cloud in a secure way.
- Adding PaaS or SaaS features like machine learning software (e.g. Microsoft Azure ML [27]).
- Adding real-time analytics engines to the clouds (e.g. [28], [29] in case of Microsoft).
- Adding tools for corporate clients to deploy IoT applications across the Cloud.
- Partnering with key players in the IoT device platform market (e.g. to support ARM's mbed OS) [35].
- Overall, getting the hyperscale cloud IoT ready (e.g. in case of Azure [30], [20]).

However, given above ongoing steps towards the IoT-ready public or hybrid Cloud, where does the evolution make use of Fog and edge cloud computing? Could real-time analytics run on Fog nodes, Cloudlets and MEC servers and therefore much closer to where many IoT devices may be located?

- b) Studying the interworking between edge cloud (in Fog, Cloudlet, ETSI MEC server) and main clouds (of Microsoft, HP, IBM, Google etc.).
- c) Studying the required interworking between edge clouds or Fogs belonging to different domains (operated by different providers).
- d) Working on how cloud-native applications (e.g. for IoT) would have to be designed to be deployable in a distributed cloud environment (of say IoT device, a fog platform or Cloudlet of provider XYZ, and the main cloud of e.g. Microsoft Azure or Amazon EC2). An interesting starting point to think this through is the use case of smart traffic light in [2].

Given where the industry is now, it wouldn't be surprising if enterprises in vertical industries (e.g. automotive, industrial) get slightly confused as they will be approached by several technology companies which offer them the vision of edge cloud computing and/or fog computing as an extension of their own product roadmaps. As one can extend Internet routers to become fog nodes (e.g. Cisco IOX [32], [33]), one can extend LTE base stations to become mobile edge compute nodes (e.g. NOKIA) and not to forget, one could extend a hyperscale cloud datacentre to reach out to the edge as well (e.g. Microsoft Azure [18]).

However, the key point is that design of a distributed cloud application for deployment across the chain of say a main datacentre, multiple fog nodes (or edge cloud nodes) and end points (the "things" in the Internet) remains a challenge.

- e) Creating ways to deploy edge/fog applications in a 3+-tier way such that the deployment process is as slick as with today's cloud offerings (e.g. Amazon EC2, Microsoft Azure etc.). For example, the paper on Fog computing [2] states: "Fog architecture exposes APIs for application development and deployment.", however it doesn't go into any details. Proprietary deployment APIs (e.g. from Cisco, NOKIA, Huawei, and others) won't do the trick I suspect and may lead to fragmentation that impedes industry adoption. However, one may distinguish two different categories of applications that could be deployed following different processes: Deploying curated applications versus deploying non-curated applications (as explained in a section above).
- f) Determining the needs and constraints of vertical industries for security and privacy in the context of distributed cloud, edge cloud and fog.
- g) Studying possible operational models: Who would operate end-to-end applications, where part of the application is hosted in a big public or hybrid cloud, other parts are hosted in edge cloud computing nodes and fog nodes and some parts are hosted on the devices? Who provides service guarantees, how do service guarantees fit together, which party is responsible for the end-to-end quality of the distributed cloud application?

The announcement of the Open Fog Consortium on 19 November 2015 [3] is very relevant, in particular as it includes Microsoft. The Fog concept needs to be seen in the context of other related industry activities as explained above. From the innovations brought to the table by the concepts of Fog, Mobile Edge Computing and Cloudlets, there is hope that breakthroughs can be achieved regarding the next generation of Cloud computing which will power much of the intelligence in the Internet of Things.

Disclaimer: Views stated, opinions expressed or analyses provided are solely those of the individual authors as private persons and shall not be considered or interpreted as views, opinions, or analyses of any of the companies they may be liaised with, employed by or do work for.

References

[1] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, Sateesh Addepalli: "Fog Computing and Its Role in the Internet of Things", Aug 17, 2012.

<http://conferences.sigcomm.org/sigcomm/2012/paper/mcc/p13.pdf>

- [2] Flavio Bonomi, Rodolfo Milito, Preethi Natarajan and Jiang Zhu: “Fog Computing: A Platform for Internet of Things and Analytics”. 2014.
<http://www.ic.unicamp.br/~bit/mo809/seminarios/Marcio-Fog/suporte/Fog%20Computing-%20A%20Platform%20for%20Internet%20of%20Things%20and%20Analytics.pdf>
- [3] Open Fog Consortium announcement 19 Nov 2015:
<http://www.openfogconsortium.org/news/internet-of-things-leaders-create-open-fog-consortium-to-help-enable-end-to-end-technology-scenarios-for-the-internet-of-things/>
- [4] Open Fog Consortium website: <http://www.openfogconsortium.org/>
- [5] ETSI ISG MEC main site (ETSI Industry Specification Group Mobile Edge Computing):
<http://www.etsi.org/technologies-clusters/technologies/mobile-edge-computing>
<https://portal.etsi.org/tb.aspx?tbid=826&SubTB=826>
- [6] ETSI ISG MEC Industry Enablement Group Terms of Reference:
https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/ToR_MEC_IEG.pdf
- [7] ETSI ISG MEC Proof of Concept Framework:
<http://www.etsi.org/technologies-clusters/technologies/mobile-edge-computing/mec-poc>
- [8] ETSI ISG MEC list of ongoing proofs of concept:
http://mecwiki.etsi.org/index.php?title=Ongoing_PoCs
- [9] ETSI ISG MEC wiki: http://mecwiki.etsi.org/index.php?title=Main_Page
- [10] ETSI NFV: <http://www.etsi.org/technologies-clusters/technologies/nfv>
- [11] Cloudlets: all about cloudlet-enabled mobile computing. <http://elijah.cs.cmu.edu/>
- [12] Cloudlets source code: <https://github.com/openedgecomputing>
- [13] Ying Gao, Wenlu Hu, Kiryong Ha, Brandon Amos, Padmanabhan Pillai, Mahadev Satyanarayanan: “Are Cloudlets Necessary?”, Oct 2015. <http://reports-archive.adm.cs.cmu.edu/anon/2015/CMU-CS-15-139.pdf>
- [14] Mahadev Satyanarayanan, Rolf Schuster, Maria Ebling, Gerhard Fettweis, Hannu Flinck, Kaustubh Joshi, and Krishan Sabnani: “An Open Ecosystem for Mobile-Cloud Convergence”, IEEE Communications Magazine March 2015,
<https://www.cs.cmu.edu/~satya/docdir/satya-cloudlet-ecosystem-2015.pdf>
- [15] Driving Cloudlet into OpenStack, OpenStack Summit Tokyo, Oct 2015.
<https://www.youtube.com/watch?v=w-gl98o3lYk>
- [16] Kiryong Ha, Mahadev Satyanarayanan: “OpenStack++ for Cloudlet Deployment”. School of Computer Science Carnegie Mellon University Pittsburgh, Aug 2015.
<http://reports-archive.adm.cs.cmu.edu/anon/2015/CMU-CS-15-123.pdf>
- [17] Victor Bahl, Microsoft, interview about micro datacentres, Sept 2015.
<http://www.networkworld.com/article/2979570/cloud-computing/microsoft-researcher-why-micro-datacenters-really-matter-to-mobiles-future.html>

- [18] Victor Bahl, Microsoft: “Emergence of micro datacentre (cloudlets/edges) for mobile computing”, 13 May 2015.
http://research.microsoft.com/en-us/um/people/bahl/Present/Bahl_mDC_Keynote_May.pdf
- [19] Guenter Klas, “About Mobile Edge Cloud Computing – An Introduction”.
<http://yucianga.info/wp-content/uploads/2015/11/15-11-08-About-Mobile-Edge-Cloud-Computing-%E2%80%93-An-Introduction-v1.pdf>
- [20] Microsoft Future Decoded 2015 Event: <https://futuredecoded.microsoft.com/>
- [21] Prof. Satya, Carnegie Mellon University, home page: <https://www.cs.cmu.edu/~satya/>
- [22] Microsoft and Rockwell Automation working on IoT enabled Cloud Computing:
<http://www.microsoft.com/en-gb/server-cloud/customer-stories/rockwell-automation.aspx>
- [23] Microsoft and Rockwell Automation on IoT:
<http://blogs.microsoft.com/iot/2014/12/04/fueling-the-oil-and-gas-industry-with-iot/>
- [24] Microsoft and Thyssen Krupp on cloud-connected elevators in IoT:
<http://www.microsoft.com/en-gb/server-cloud/customer-stories/thyssen-krupp-elevator.aspx>
- [25] Microsoft announces plan to offer cloud services from the UK, 10 Nov 2015:
<http://news.microsoft.com/en-gb/2015/11/10/microsoft-announces-plans-to-offer-cloud-services-from-the-uk/>
- [26] Microsoft plan to open two German-only Azure Data Centre Regions.
<http://www.techweekeurope.co.uk/cloud/microsoft-azure-german-data-centre-180433>
- [27] Azure Machine Learning (Azure ML): <https://azure.microsoft.com/en-gb/services/machine-learning/>
- [28] Azure Stream Analytics: for real-time analytics. <https://azure.microsoft.com/en-gb/services/stream-analytics/>
- [29] Cortana Analytics Suite: <http://www.microsoft.com/en-gb/server-cloud/cortana-analytics-suite/what-is-cortana-analytics.aspx>
- [30] Microsoft Azure IoT Suite: <http://www.microsoft.com/en-gb/server-cloud/internet-of-things/azure-iot-suite.aspx>
- [31] Adrian Bridgewater, Forbes, 11 Nov 2015. “Microsoft Azure Cloud Evolves For Intelligent Machine Learning”. <http://www.forbes.com/sites/adrianbridgewater/2015/11/11/microsoft-azure-cloud-evolves-for-intelligent-machine-learning/>
- [32] Cisco IOx: <https://developer.cisco.com/site/iox/>
- [33] Cisco IOx presentation, Jeremie Garnier, 4 Nov 2015:
<http://www.slideshare.net/CiscoDevNet/iox-techintroforparishackathon>
- [34] OpenEdgeComputing.org: <http://openedgecomputing.org/about-oec.html>
- [35] ARM mbed – Microsoft Azure cloud integration:
<http://blog.mbed.com/post/130148704027/cloud-integration-gets-simpler>