Contributors
Mahadev Satyanarayanan, Carnegie Group Professor of Computer Science, Carnegie Mellon University
Dimitra Simeonidou, CTO, Bristol is Open
Rolf Schuster, Director, Open Edge Computing (OEC) Consortium

Edge Computing for Network Operators

OPPORTUNITIES, BUSINESS MODELS AND PARTNERSHIPS

Alexander Cherry

9 March 2016

Edge Computing for Network Operators: Opportunities, Business Models and Partnerships

Cloud 2.0, or Why We Need Edge Computing

You might be forgiven, in this world of everything-as-a-service and amid the ever more grandiose predictions as to the eventual size of the IoT, for thinking that the golden age of telecoms is forever at an end. Gone are the days of the crackling internet dial-up tone. Gone is the SMS message. The cloud-computing revolution has placed seemingly unlimited computing resources at our fingertips, and in our infatuation with the cornucopia of new services that this has enabled, both B2B and B2C, we have largely forgotten the physical reality of the internet. Most internet subscribers treat telco access networks in much the same way as they treat the roads they drive down on their way to the shops, to the cinema, to their favourite restaurant. Money is in retail, money is in film, money is in catering; money is not in roads.

However, it is not the age of telecoms that is nearing its end; rather it is the end of the illusion of an omnipotent, ubiquitous, boundlessly scalable cloud. For as we have multiplied exciting new over-the-top (OTT) services hosted centrally, so too have we multiplied devices and sensors all around the physical world that are eager to consume these services – witness the oft-quoted Cisco prediction of 50 billion connections by 2020, of which, according to TechCrunch, 6.1 billion are likely to be data-hungry smartphones[i]. This explosion of supply in central clouds and demand at the very edges of the network is placing unprecedented strain on the infrastructure in between. Firstly, bandwidth is limited, no more so than in the mobile portion of the network, so there is a limit to how much information can be relayed at once; in the case of video for instance, streams must be squashed (leading to a loss of picture quality) or stretched thin (leading to buffering's dreaded spinning-wheel of death). Secondly, the greater the distance between the end point and the cloud, the slower will be an application's response, making strictly real-time applications impossible or suboptimal within the current cloud paradigm.

Our desire for ever greater bandwidth and lower latency has therefore set us on a collision course with the physical limitations of internet infrastructure: with the miles of copper and fibre cables hidden from view, with the routers and switches, with the mobile cell towers. Having for too long been the literal conduits of other people's innovations, telcos now have a central role to play in overcoming the bandwidth and latency constraints inherent in today's centralised cloud paradigm, by leveraging their networks' distributed footprint to host applications as close as possible to subscribers or, in industry parlance, 'at the edge'. With fixed networks extending into businesses and homes, and mobile networks providing near-ubiquitous connectivity to the habitable world, telcos have a unique and compelling opportunity to implement edge computing in their networks, not just meeting the ever-rising

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

demand for internet data but enabling a new breed of high-bandwidth, low-latency applications that, crucially, they can monetise. Network infrastructure, so long the ugly duckling of the internet age, is on the verge of becoming a swan.

In order to learn more about the opportunity that edge computing represents for network operators, the applications and business models it enables, and the partnership models under which these can be developed quickly and profitably, we recently spoke to *Mahadev Satyanarayanan (Satya),* Carnegie Group Professor of Computer Science at **Carnegie Mellon University** (CMU), *Dimitra Simeonidou*, CTO at **Bristol is Open**, and *Dr Rolf Schuster,* newly appointed Director of the **Open Edge Computing (OEC) Consortium.**

Why Network Operators Need to Take Notice of Edge Computing

'Effectively, by providing edge computing, operators have a unique, one-time opportunity to move up the value chain, where significantly higher profit margins are possible'

Satya Mahadev, Carnegie Group Professor of Computer Science at Carnegie Mellon University (CMU)

Operators have, since the dawn of the internet 20 years ago, continually ceded higher-margin services to OTT players and seen their core business – connectivity – steadily commoditised.

'Today it is the case that operators, much to their disgust unfortunately, are essentially bit pipes,' comments Satya, 'all of the value lies in the cloud service, in Facebook, in Google, in other cloud service providers.'

While this does mean more data traffic, operators cannot monetise this as effectively as, for example, voice and SMS, their old bread and butter.

Rolf, recently appointed Director of the OEC Consortium after nearly a decade at Vodafone, sees the current trend in telecoms in a similar light: 'The problem is that there is rapidly growing demand for data but the prices are not growing dramatically, they stay flat, you have more traffic but you don't get more pay.'

For all their forays into higher-margin OTT services of their own – video streaming, messaging apps and cloud hosting to name a few – these attempts by operators to add new revenue streams to their core connectivity business have remained largely fruitless. You might say that operators have been trying to beat the OTTs at the OTT game and have failed. It is now time for them to play their own game though, and to play to their strengths.

As Satya sees it, operators' unique value proposition consists in their distributed footprint, which extends all the way to the subscriber edge: 'I think the compelling argument for why the operator is uniquely positioned is precisely because of this.

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

They own the real estate upon which your packet first meets the infrastructure, and widely structuring and using that real estate is going to be crucial.'

The reason this real estate at the edge of the network is so valuable is because it is a prime location for hosting any bandwidth- or latency-sensitive application. 'In real estate it is a cliché that location, location and location are the three most important attributes for value,' Satya explains. 'The unique new opportunities for operators,' he continues, 'lie in enabling applications that can sit and benefit from this proximity to the edge.'

Rolf frames the edge-computing proposition in similar terms: 'There is value in bringing computing and storage resources closer to the customer, and who is closest to the customer? The operator. They have the fixed line side, they have the mobile side as well, so if there is demand for storage and compute resources close to the customer then the operators are well-positioned to utilise that and offer that.'

So what is the big problem with the central cloud when we're talking about bandwidth- and latency-sensitive applications? As it stands, round-trip communication between edge devices (the consumers of applications) and central clouds (the hosts of applications) takes on average 70 – 100ms, though this figure varies massively from packet to packet depending on conditions across the internet and does not take into account any processing time for the application itself. This may not seem like a lot but there are certain types of application that it can impact significantly. Take for instance Augmented and Virtual Reality – a market TechCrunch reckons will be worth $150 billion by 2020[ii].

Imagine you were wearing an Augmented Reality headset that overlaid information on the world around you. It would only take you a split second to turn your head by 90° or 180°, exposing a whole new visual field, and you would want your overlay to update itself as you turned; studies have in fact shown that AR and VR, once lag or jitter becomes perceptible, will generally cause chronic motion sickness. Another example: in just a quarter of a second a car travelling at 100km/hour moves nearly 7m, making autonomous-driving applications difficult to run out of the cloud, as the lay of the road and relative positions of the cars may, by the time an answer arrives, be quite different from how they were when the question was asked.

Summary of Latencies Achieved with an Unnamed Major Cloud Provider

Data Center	Average Latency
West US	115ms
South Central US	131ms
East US	155ms
North Central US	171ms
North Europe	222ms
West Europe	223ms

Source: 'Emergence of Micro-Datacentres (Cloudlets/Edges) for Mobile Computing' (Microsoft Research, 2015)

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

If an application is hosted not in a central cloud but in a local cloudlet deep inside an operator network, close to where it is needed, the latency at which it runs can be radically improved, with round-trip times falling to ~20ms and lower. While the economies of scale for cloudlets are less than for massive centralised clouds, the hosting service they offer is potentially a lot more valuable and therefore worth more money – because they radically enhance what we can do already and outright enable a whole host of so-called 'tactile internet' applications, including (mobile) AR/VR experiences and real-time IoT/M2M, where a lot of money stands to be made over the coming years.

As Satya puts it, 'Effectively, by providing edge computing, operators have a unique, one-time opportunity to move up the value chain, where significantly higher profit margins are possible.' He goes on: 'If you can provide exactly what cloud hosting providers are providing but closer to the edge, and at lower latency and possibly improved scalability, then in fact you can charge premium pricing for those services.' It is time for operators to move away from being mere sellers of bandwidth, towards higher-value services.

The business drivers seem pretty clear, but what does edge computing mean in practical terms? Regardless of whether we're talking **Mobile Edge Computing** (MEC), which is currently being standardised by ETSI and developed through a series of proofs of concept (POCs), **Open Edge Computing** (OEC) or **Fog Computing**, it involves the establishment of small clouds, often termed **cloudlets**, at strategic nodes in the network.

'A beautiful way to think about cloudlets is that they do for general purpose computation what CDN did for data delivery,' says Satya, whose 2009 paper *The Case for VM-Based Cloudlets in Mobile Computing* brought the term into the mainstream. 'It's compute on demand, close to the edge of the network.'

For telcos this potentially means adding cloud-like capacity anywhere from a mobile base station or an aggregation point to a Central Office (CO). 'One of the areas that operators will be sorting out as they pilot edge computing is precisely where to locate the edge that is going to host all these compute capabilities,' clarifies Satya. 'My suspicion is that there will be some applications that are so extremely latency-sensitive, sub-millisecond, where you need to be really at the cell tower. However, my expectation is that it's mostly going to be deeper, so that, in the core network, maybe over a multi-cell-tower coverage area, you have a single cloudlet.'

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

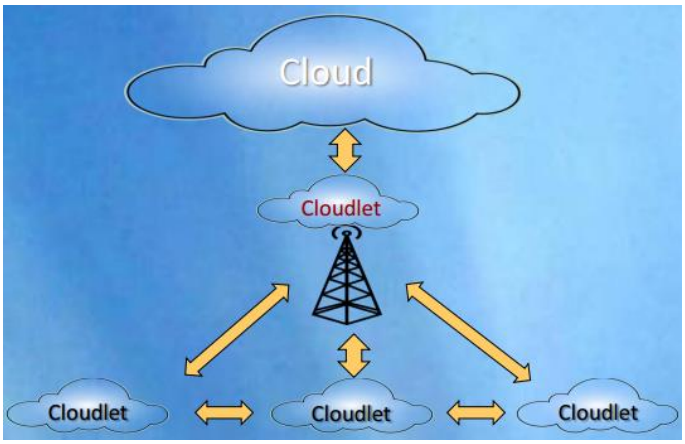

Source: 'On the Advanced 5G Network Infrastructure for the Future Internet' (Huawei, 2013)

Historically, the candidate host sites for cloudlets – base stations, aggregation points and Central Offices – have been host only to specialised signal-processing, so cloudlets might seem like a somewhat unnatural or unwieldy bolt-on. However, regardless of where operators stand on edge computing, this trend is already well underway in their networks in the form of **Network Functions Virtualisation** (NFV) and **Software-Defined Networking** (SDN). It is but a small step from hosting Virtual Network Functions on commodity hardware distributed within the network to hosting *any* virtual function on commodity hardware distributed within the network – and this is essentially all a cloudlet is.

As Rolf says, 'If operators have converged their networks into an NFV infrastructure, then at the end of the day you have virtualised cloud services anyway across all network components, and then calling one NFV node, or part of it, an edge server is not really a new investment anymore.'

With the continued maturation of NFV and SDN, as of 2016 the facilitating technologies for edge computing in telco networks are falling into place at the same time as the business drivers are coming into sharper focus: the network edge stands to be one of the key strategic battlegrounds where the future of internet services is decided.

Show me the Money: Applications and Business Models at the Edge

The current centralised cloud paradigm falls down primarily in terms of bandwidth and latency, and it is in these terms that the new generation of edge-based applications can best be categorised. Content/video delivery is perhaps the most mature use case, and the closest to operators' existing business. The 'tactile internet' on the other hand encompasses a wide range of emerging low-latency applications (the idea is that the internet should respond immediately to our inputs, as if to our touch).

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

Edge Computing Application Type	Bandwidth-Heavy	Latency-Dependent
Content and Video Delivery	High	Low
Tactile Internet I (AR, VR)	High	High
Tactile Internet II (IoT and M2M)	Low	High

We will be exploring in turn each 'family' of edge computing applications in a series of follow-up whitepapers, though a few pointers here will give an idea of the size of the market in each case:

1) **Content and Video Delivery:** while consumers don't mind having to wait a second or so for their favourite programme to load, what really undermines their viewing experience is buffering, quality degradation and jitter, something particularly prevalent in mobile networks where links are 'longer' and 'thinner' than at the heart of the internet. The application of edge computing to video is generally not a latency play, it is a bandwidth play – and will only get more important with over 80% of the world's internet traffic slated to be video by 2019. [iii]

2) **AR and VR:** these represent two of this decade's biggest emerging markets, with an estimated collective worth of $150 billion by 2020 and applications across a multitude of verticals, including tourism, retail, healthcare and facilities maintenance. Both require extremely low latency and no jitter if they are to offer a tactile response to our head movements, as well as a lot of bandwidth, making edge computing the perfect enabler, especially in mobile scenarios.

3) **IoT and M2M:** Countless IoT applications, from industrial automation and autonomous driving to remote/robotic surgery, are completely unthinkable with even so much as the smallest hint of latency. There is barely a need to repeat Cisco's infamous prediction of 50 billion connected devices by 2020 to see the size of the opportunity in this field as well.

Satya believes that operators' business case for deploying edge computing in their networks will naturally be driven by a variety of new use cases – each representing a slightly different opportunity for operators, and some as yet totally unthought-of. The exact business model by which operators will monetise cloudlets within their networks also needs to be worked out in full, though several models are currently being discussed. Rolf outlines three potential candidates:

- **Free Model:** applications are allowed into the telco network free of charge, and the charge is passed over to the consumer in some way (e.g. a 'premium' internet subscription).

- **Retail Model:** the telco sells cloudlet capacity to application developers and manages these client relationships directly. This would effectively see telcos

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

enter into competition with centralised cloud providers within certain latency- and bandwidth-sensitive horizontals.

- **Wholesale Model:** the telco sells cloudlet capacity to an established cloud provider, who then resells that capacity to its own roster of application-developer clients, thereby complementing its central hosting offer with support for low-latency or bandwidth-sensitive app components.

The retail and wholesale models have so far gained the greatest credence, promising as they do a direct source of new revenues for network operators. To illustrate how the retail model could work in more concrete terms, Satya gives the example of an operator establishing a cloudlet in close proximity to a shopping mall. An AR company could host its application in this cloudlet in order to outright enable, or at least significantly enhance, its AR offering to store owners, whose customers would, in turn, enjoy richer AR experiences and finally be in a position answer that perennial question: what will this dress look like on me? The bottom line is more satisfying shopping experiences for consumers, more money for store owners, more money for AR companies and more money for the host network operator, whose cloudlet makes everything possible in the first place.

The choice between retail and wholesale models will depend on a host of factors specific to each application, and in all likelihood these two approaches will complement each other. In verticals where telcos already have extensive business relationships and sales capacity (e.g. media content distribution), it is easy to see them adding cloudlet capacity into the mix of products they already sell.

However, there are plenty of candidate industries for edge computing where these relationships simply don't exist, and telcos would be required to establish a wholly new sales effort in order to fully develop the opportunity. In these cases, we can see the value of the wholesale model. Rolf believes that having a cloud service provider between the operator and the application developers is important here because the cloud service providers already have thousands of customers that they can bring to the table.

'The operators may struggle to build their sales channels into that space and build reliable services on a global scale,' cautions Rolf, 'and don't forget that some of these application providers want to have global reach, so we have to avoid the situation where they have to deal with 150 different operators.'

With the wholesale edge-computing model, what we have is a potential win-win situation. Telcos can focus their energy on developing and managing the infrastructure, while cloud service providers can add an extra dimension to the hosting services they already offer their clients. Rolf sees this as a natural extension of cloud service providers' moves over the past few years towards ever more local Points of Presence (POPs) or 'edge locations', and believes that by extending their footprint deeper into operator networks they will be able to differentiate their hosting service and justify higher prices to their clients.

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

Developing the Edge Opportunity: Third Parties, Partnerships and POCs

'There are hungry people out there, hungry entrepreneurs, who are hungry to create new applications, and there are new devices'

Satya Mahadev, Carnegie Group Professor of Computer Science at Carnegie Mellon University (CMU)

With limited R&D budgets and so many possible avenues to explore, it can be difficult for operators to know where to start with edge computing. Whatever they do, a fresh approach is going to be necessary. Telcos' core connectivity business is relatively commoditised, and this is reflected in the business model that underpins it, which for a long time has been all about scale and consistent Quality of Service (QoS). There is scarcely a need to elaborate a business case for on-going investment in infrastructure – all new capacity will be eagerly taken up by consumers. If you build it, they will come, so the key is to build it as cost-effectively as possible. This has generally been achieved by realising economies of scale on network investments, namely by building massively and uniformly in drawn-out CAPEX cycles.

While some use cases, like transparent caching of media content, can reduce on-going network cost, edge computing is fundamentally a new-revenue play. If the new services it unlocks generate higher margins than the steadily commoditised connectivity business, then they need not be subject to the same tyranny of scale; a one-off local deployment of edge computing is economically viable in a way that a one-off local network build-out might not be. Besides, edge computing is still such an unknown quantity – with a vast range of different use cases and potential deployment topologies – that there is not yet a 'template' that could serve as the basis for a scaled deployment. Edge computing requires a new mentality: operators need to experiment with piecemeal, bespoke deployments addressing particular use cases rather than waiting until some kind of master plan is in place for a mass deployment. Rolf believes they should start with local or regional roll-outs, and this sentiment is echoed by Satya.

'If you only roll out edge computing when there's enough demand, a very high demand for it everywhere,' Satya warns, 'it's going to take a long time and the opportunity is going to pass.'

Both Rolf and Satya consider direct engagement with third-party industry verticals to be a vital step towards creating a true edge-computing ecosystem. 'The biggest problem is the application demand,' Rolf explains, 'and that's a typical chicken-and-egg problem that every technology enabler has.'

Application developers creating applications with centrally located clouds in mind are understandably reluctant to invest in apps for as yet non-existent infrastructure; likewise, telecom operators are reluctant to invest in infrastructure for as yet non-existent apps. The key difficulty telcos face, and will continue to face, in developing POCs for edge deployments is that a lot of edge use cases lie outside their traditional

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

business areas. A similar consideration applies to third-party developers, many of whom are familiar with the cloud world but less so with telecoms.

As Rolf puts it, 'You have to bring 3 or 4 different industries together, definitely the telecom and the IT industry, and that never is easy and they work differently and on different timescales.'

Fortune favours the brave, and operators need to get out of their comfort zones in terms of the communities they interact with, and engage directly with third parties in those industries where edge computing can make a real difference.

'I think they should hold events and reach out directly with them,' proposes Satya, before adding, 'being proactive and incentivising the developer community to the infrastructure of the future – that is in the interest of the operators and they should do it now.'

Indeed, we are seeing the emergence of tripartite arrangements in which operators, application developers and network vendors work in unison to develop the business case, and business model, around particular edge use cases. Consortia, government initiatives and academia therefore have an important role to play, insofar as they can incorporate perspectives from both sides of the table and act as a kind of rallying point or middle ground for multiple industry partners.

In order to learn more about the value of such intermediaries, we spoke to Dimitra Simeonidou, CTO at Bristol is Open, a government- and industry-sponsored testbed that has created an advanced citywide network in Bristol, with real users, in a bid to become 'the first SDN city in the world'. Bristol is Open is currently testing the value of both cloud and edge computing for specific smart-city applications, including CCTV video analytics, congestion management, driverless cars and robotics. Some of these applications are well served by high-firepower central cloud models. Others though require extremely low latencies, which makes them obvious candidates for edge computing.

'We are moving very much to seeing deployed probably the first fully operational 5G testbed in the UK, and you know that for 5G consideration one of the big KPIs is latency,' says Dimitra. 'Therefore moving your computational resource very close to where things are going, very close to your interface, becomes increasingly important.'

Bristol is Open have an established partnership with InterDigital, with whom they have developed the Flexible IP-Based Services (FLIPS) concept, which uses edge computing for advanced media delivery and was recently validated as an official POC by ETSI's MEC Industry Specification Group. Another partnership with network vendor NEC was announced more recently with a focus on leveraging distributed computing for IoT applications in the smart city.

'The reason we are working with big corporations, be they vendors or operators or others, is because they are looking to use our infrastructure to try and experience new services and new technologies,' Dimitra explains, 'we are going to give them an

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

environment to try out what is the next generation of cloud infrastructure for specific applications.'

As Dimitra sees it, the key advantage that an initiative like Bristol is Open offers to network operators is 'the possibility to try new things with very low entry because the investment has already happened in order to put this capability together'. It simplifies outreach to specific industry verticals, to independent R&D units, to academia and to end consumers, allowing telcos to focus on the commercialisation aspect of these new technologies. Dimitra concludes: 'This is exactly where we are making an impact, because we have a testbed which is out in the wild and deployed with real users.'

Satya also emphasises the value of gaining real-world experience in this way, adding that, 'people need to experiment to understand what works, what doesn't, and to understand different business models and what makes money.' This will ultimately build the confidence that the industry needs to get investment flowing, both on the operator and on the application-developer side.

'There are hungry people out there,' says Satya bullishly, 'hungry entrepreneurs, who are hungry to create new applications, and there are new devices.'

While he recognises that it is still early days, he is adamant that telcos must not sit back and wait for the opportunity to develop by itself. If they do, he worries that third-party companies could come along instead and position cloudlets very close to the edge of the network but outside the operators' control. And if this happens, it won't just be the operators that are poorer – tomorrow's applications will be slower and less available for consumers and businesses alike than if they run at the operator edge.

A Journey through the Fog

'The best way to predict the future is to invent it' Alan Kay

As with any new technology, there will be a tipping point after which edge computing becomes self-sustaining; the more widely available cloudlets are, the more applications will get written to leverage them, increasing demand further and initiating a positive feedback cycle. In the meantime, it is vital that network operators take a proactive role in driving things forwards.

Satya emphasises the value of operators working with one another: 'We're talking about creating a whole new market, and I think the operators have to recognise that their interests are actually all aligned, rather than saying, that particular application is created by Verizon, here's another one that's going to be created by AT&T… all these operators, and other telco players, they all have an incentive, a cause, to accelerate the creation of new applications that can leverage the edge.'

Telcos are an important part of this new emerging value chain but, as Satya's statement implies, there are other players that will make money out of edge computing. Rolf agrees, noting that 'the whole cloudlet concept has to be bearer-agnostic, so I would even include fixed line, I would include Wi-Fi, I would include

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

ZigBee, whatever you have, Bluetooth.' It is also likely that systems integrators and (Mobile) Virtual Network Operators will build their own value-added solutions on top of operators' underlying edge offering.

Another man's cloudlet should not be seen as an enemy though, and the industry as a whole must do whatever it can to encourage interoperability. Application developers are unlikely to care for instance whether they are delivering their app over an operator-managed LTE edge or over a Wi-Fi edge, and will push for availability of the largest possible footprint of computing resources. Satya expands upon this future vision: 'The application creator would preferably like to see a system whereby an application that he creates uses Wi-Fi-based cloudlets if that's where you're using it, telco cloudlets if you happen to be on 4G; maybe an automobile cloudlet via Bluetooth or Wi-Fi if you happen to be in a car; and so on. It's important to recognise this particular point: that there are going to be many forms of cloudlets.'

A linked concern to portability *between* different operators' cloudlets is portability from edge to centre, a distinction that will surely become more and more arbitrary as we move towards ever more virtualised networks under NFV. For the time being, Rolf believes edge computing should be kept separate from NFV, though in close touch, so that it can develop faster. Ultimately he does see convergence between the two; in the future, application developers will not choose to leverage 'edge' capabilities as opposed to 'core' or 'cloud' capabilities, they will simply want their app to run at whichever point in the dynamic, intelligent network represents the sweet spot for performance and cost.

This application- rather than technology-driven point of view is echoed by Dimitra. 'The whole city is a datacentre,' as she says of her work at Bristol is Open – 'you should be able to manage all this environment of resources as a single cloud'. She continues: 'We are driving full integration between the network and the IT resource wherever the IT resource is sitting, and really what we are looking at for the future are architectures and protocols whereby we can actually support full disaggregated environments.'

A fragmented operating environment – whether this fragmentation is vertical, from cloudlet stack to cloudlet stack, or horizontal, from central cloud to edge cloudlets – will act as a drag on the industry, hurting everyone in the value chain, from application developers through telcos and other edge owners right down to consumers. ETSI's MEC ISG has set itself the task of standardising APIs across 4G/LTE cloudlets, and this will help to eliminate walled gardens in the Mobile Network Operator (MNO) community and to grow the accessible market for application developers – but operators should look further still and seek to include other types of network and other connectivity formats in a broader edge/cloudlet vision. This is certainly what Rolf and Satya hope to achieve through the action of the OEC Consortium, closely aligned to ETSI's MEC ISG but seeking to standardise a single API for all cloudlets. The bigger the overall edge/cloudlet pie, they reason, the bigger operators' slice will be, whatever this slice ultimately contains.

Standards bodies like ETSI's MEC Industry Specification Group, the Open Edge Computing Consortium and the recently established OpenFog Consortium each represent slightly different flavours of edge computing, and approach the

White paper produced in association with **Network Edge Europe 2016** (6-7 June, London). Find out more here: http://openmobilemedia.com/networkedgeeurope/

distributed-cloud challenge from different starting points. These industry efforts are valuable stages towards a convergence that we believe will be bigger than any single one of them and encompass broader developments in cloud computing, NFV/SDN and the still hazy notion of 5G.

Whatever form it ultimately takes and whatever term the industry ultimately settles upon to describe it, this much is crystal-clear already: edge computing can offer network operators that long-sought role back at the heart of the internet revolution.

This white paper has been produced in association with Network Edge Europe, which takes place on 6-7 June at the Royal Garden Hotel, London. The event is free to attend for senior-level telcos and features over 20 speakers from across the application developer, telco and enterprise sectors.

For more information visit http://openmobilemedia.com/networkedgeeurope/

[i] http://techcrunch.com/2015/06/02/6-1b-smartphone-users-globally-by-2020-overtaking-basic-fixed-phone-subscriptions/

[ii] http://techcrunch.com/2015/04/06/augmented-and-virtual-reality-to-hit-150-billion-by-2020/

[iii] www.reelseo.com//2019-internet-video-traffic/