Technology Institute

Mobile Innovations Forecast

Virtual context: Connecting two worlds

Raman Chitkara, Global Technology Industry Leader

Pierre-Alain Sur, Global Communications Industry Leader

Contextual services will accelerate the shift to programmable networks

When people upload mobile photos or video to a social network, back up documents and media files to the cloud or sync their data from a mobile device, they typically cross multiple network boundaries and data centres. In these scenarios, latency, jitter and packet loss affect user perceptions of network quality just as pops, crackles or dropped calls impact perceptions of mobile voice quality.

These data quality metrics have always been important, but they take on greater significance as people use increasingly sophisticated, context-aware mobile devices and services. As Qualcomm CEO Dr. Paul Jacobs has noted, wherever users physically travel, they will be at the centre of a personal cloud of nearby devices, apps, information and contextual choices. "We're working on a vision we call the Digital Sixth Sense, this idea that the world around us will be connected, and the phone and devices we carry will allow us to essentially blur physical space and cyberspace," he recently told PwC in a interview for their 17th Annual CEO Survey.[1] Now imagine a sensor-connected, visual media-centric world in which the end user's personal cloud of devices, information and applications continuously interacts and exchanges data with numerous service provider clouds depending on various triggers, such as geofences, in the physical world.

1 http://www.pwc.com/us/en/ceo-survey-us/2014/assets/dr-paul-e-jacobs.pdf

Synopsis

This is the third article in Mobile Innovations Forecast Phase II: New Technological Capabilities. Here's a roadmap of the series:

- The Introduction argues that the dominant drivers of mobile innovation to 2018 will revolve around capturing and modeling the contextual situation of mobile users, and will transform the mobile device into an intelligent digital assistant.
- The second article examines how device and environmental sensors interact to capture information to model the user's physical context.
- **This third article explores how communications networks will enable interaction of the user's physical context data with information and applications in the cloud to create the virtual context layer.**
- The fourth article will describe the modeling, intelligence and analytic engines, mainly in the cloud, that will enable the mobile device to become as intimate as a personal assistant, if you allow it.
- A concluding article will highlight the most significant new capabilities driving smart devices towards true digital companionship, setting the stage for new use cases and business models to follow. ■

This article focuses on information exchange between the digital and physical worlds that creates the virtual context of the end user. It is concerned with how the networking requirements of contextually aware and intelligent services will enable the dynamic environment described by Jacobs. It builds upon previous work in this series that focused on enabling technologies in Phase I,[2] and the physical context of users identified in the preceding article.[3] A fourth article will examine how cloud-based analytic and predictive engines organise and make sense of the massive data flows handled by mobile networks to create a seamless, contextualised user experience.

A killer ecosystem rather than a killer app

The virtual context of a user is created by the layer of telecommunications technologies and services that connects situational data captured by the mobile device with data, analytics and applications in the cloud. Telemetry from device sensors, location/object beacons, user-generated media and physical motion is uploaded into networks whilst interactive maps, augmented reality visualisations, streaming media and other contextually relevant information is downloaded from networks for display on a user's mobile device.

Virtual context is a dynamic environment in which a user's personal cloud of devices, data, preferences, applications and social connections constantly updates and adjusts its behaviour based on its interactions with the physical world plus information and applications accessed from networks. "You want to integrate a specific location and a specific activity you're doing with the digital information you have from your social network, your email communication, the notes you take on the device and so forth," says Oliver Brdiczka, director of contextual intelligence research at PARC, a Xerox company.

Making such an experience robust and seamless to the user requires significant network bandwidth. However, raw capacity alone is not sufficient to enable diverse mass-market contextual services. Networks must become flexible enough to not only provide various levels of quality of service, but also handle security and other back-end services such as billing for end users and third-party service providers. For that to occur, communications networks must become more than just large data pipes. They must enable the digital equivalent of smart logistics for telecom and non-telecom service providers for whom the network is the front door to the customer.

For example, a media company might want the communications network to emphasise low-latency for streaming files with a thin path coming back. At the same time, a gaming company looks for low-latency for both the uplink and downlink. Healthcare service providers need network performance plus security that is technically compliant with regulatory objectives. Search engines will start to receive images uploaded by the user in addition to key words before returning a result.

Practically speaking, the need for flexible, robust digital logistics will push cloud computing principles and technologies deeper into the design and operation of communications networks. Both communications infrastructure and service providers are virtualising telecommunications in much the same fashion as the IT industry virtualised computer processing and storage. This activity is accelerating the creation of a new capability—the programmable network.

The rise of programmable networks

Programmable networks are those in which software directs the flow of data and the behaviour of network elements in a manner that is largely independent of physical hardware. It is a process not dissimilar from the unbundling of software from mainframe computers. The result is that administrators are able to re-programme a communications infrastructure instead of re-build an infrastructure whenever they want to change or modify the services that run on top of it.

Contextually intelligent mobile technology and services are likely to accelerate adoption of programmable networks. As more powerful sensors and mobile devices capture environmental data that must be correlated across multiple data centres with stored user data, then analysed and returned to the user in real-time as a personalised, contextually relevant suggestion or action, the network infrastructure must become robust enough to handle communications traffic that exhibits both high volume and high complexity. [See sidebar on page 5]

Communications Services Providers (CSPs) are attacking the capacity and complexity problem along two broad fronts. To address high volume, they are deploying new radio interfaces and all-IP infrastructure based largely around the Long Term Evolution (LTE) standard. In addition, CSPs are pushing capacity deep indoors by nesting multiple large cell and small cell networks within one another—called heterogeneous networks (HetNets)—to load or offload user traffic onto the most optimal network infrastructure at any one time.

But handling high volume mobile data traffic is only one side of the coin. The complexity challenge for CSPs is equally acute as the architectural merging of the Internet and telecommunications

2 Wrapping up Phase I: New data bolsters the general direction of innovation over the next five years

3 Sensing and making sense: Device and environment underpin contextually aware services

becomes increasingly mature. The transition to an all-IP future through LTE and similar technologies also enables new core architectures, such as Software Defined Networks (SDNs) and Network Function Virtualisation (NFV). These approaches abstract physical switches, routers and other components into a single virtual network layer that can be managed centrally. This allows CSPs to partition or 'slice' the same network infrastructure into application- or industry-specific functions. An SDN/NFV approach to architecture also enables communications networks to behave as platforms for service innovation at the CSP level and by third parties.

Whether voice or data, most mobile user sessions begin indoors or have some aspect of them that happen within a physical structure such as a building or an automobile. Not surprisingly, better and more powerful indoor capacity is an imperative for CSPs.

The end result is a telecommunications network that acts increasingly like a computing cloud.

Ubiquitous capacity

Whether voice or data, most mobile user sessions begin indoors or have some aspect of them that happen within a physical structure such as a building or an automobile. Moreover, people are bringing more sophisticated mobile devices and sophisticated expectations to their work, school and play. Not surprisingly, better and more powerful indoor mobile data capacity is an imperative for CSPs.

To meet consumer expectations of a seamless mobile experience, CSPs have started to assess both network performance and the quality of user experience as a function of app coverage more than voice coverage. App coverage measures the likelihood that a network will deliver sufficient performance to run a particular application at a quality level acceptable to the user. From the perspective of a CSP, app coverage is how well the mobile data network performs at the geographic edge of a cell. Partly in response to the need for better coverage for mobile apps, CSPs are installing new radio interfaces like 4th Generation (4G) LTE for higher capacity, and employing small-cell technologies to complement indoor and outdoor coverage.

In principle, LTE offers three broad benefits for dealing with mixed voice and data traffic. The first is a 10X increase in mobile data rates via Orthogonal Frequency Division Multiplexing (OFDM) compared to the 3G data transmission technology known as High Speed Packet Access (HSPA). Maximum mobile data rates according to the LTE standard include 300Mb/sec on the downlink and 75Mb/sec on the uplink. This increased speed also comes with increased spectral efficiency, up to 3X more capacity per bearer channel than a typical circuit-switched 3G network. Finally, LTE achieves one-fourth the latency (data packet transfer time from sender to receiver) of comparable 3G networks.

Along with adopting LTE in the macro-network, CSPs are expanding their use of HetNets for addressing high volume mobile data traffic in outdoor and indoor environments. HetNets attack the indoor/outdoor coverage problem along three broad fronts: the macro-cellular network provides wide-area broadband coverage from the outside; a dense mesh of enterprise, metro and small-cell technologies are linked inside for high-traffic areas such as airports or offices; and network intelligence steers traffic between macro and small-cell networks according to current demand to give the user a consistent indoor/outdoor experience [see graphic on page 4].

Two broad categories of small-cell technology underpin much of the current and future strategies for indoor radio coverage. Both are based on distributing antennas whilst concentrating base stations and backhaul. The current indoor coverage model is Distributed Antenna Systems (DAS). DAS architectures split the transmitted power of a single high-powered indoor antenna into a group of low-powered antennas over the same area. The placement of extra antenna elements, albeit at lower power, helps network designers work around

continues on page 6

Figure 1: Anywhere a crowd gathers, so do devices and applications

Heterogeneous networks provide data and voice capacity for complex, high volume traffic indoors and outside.

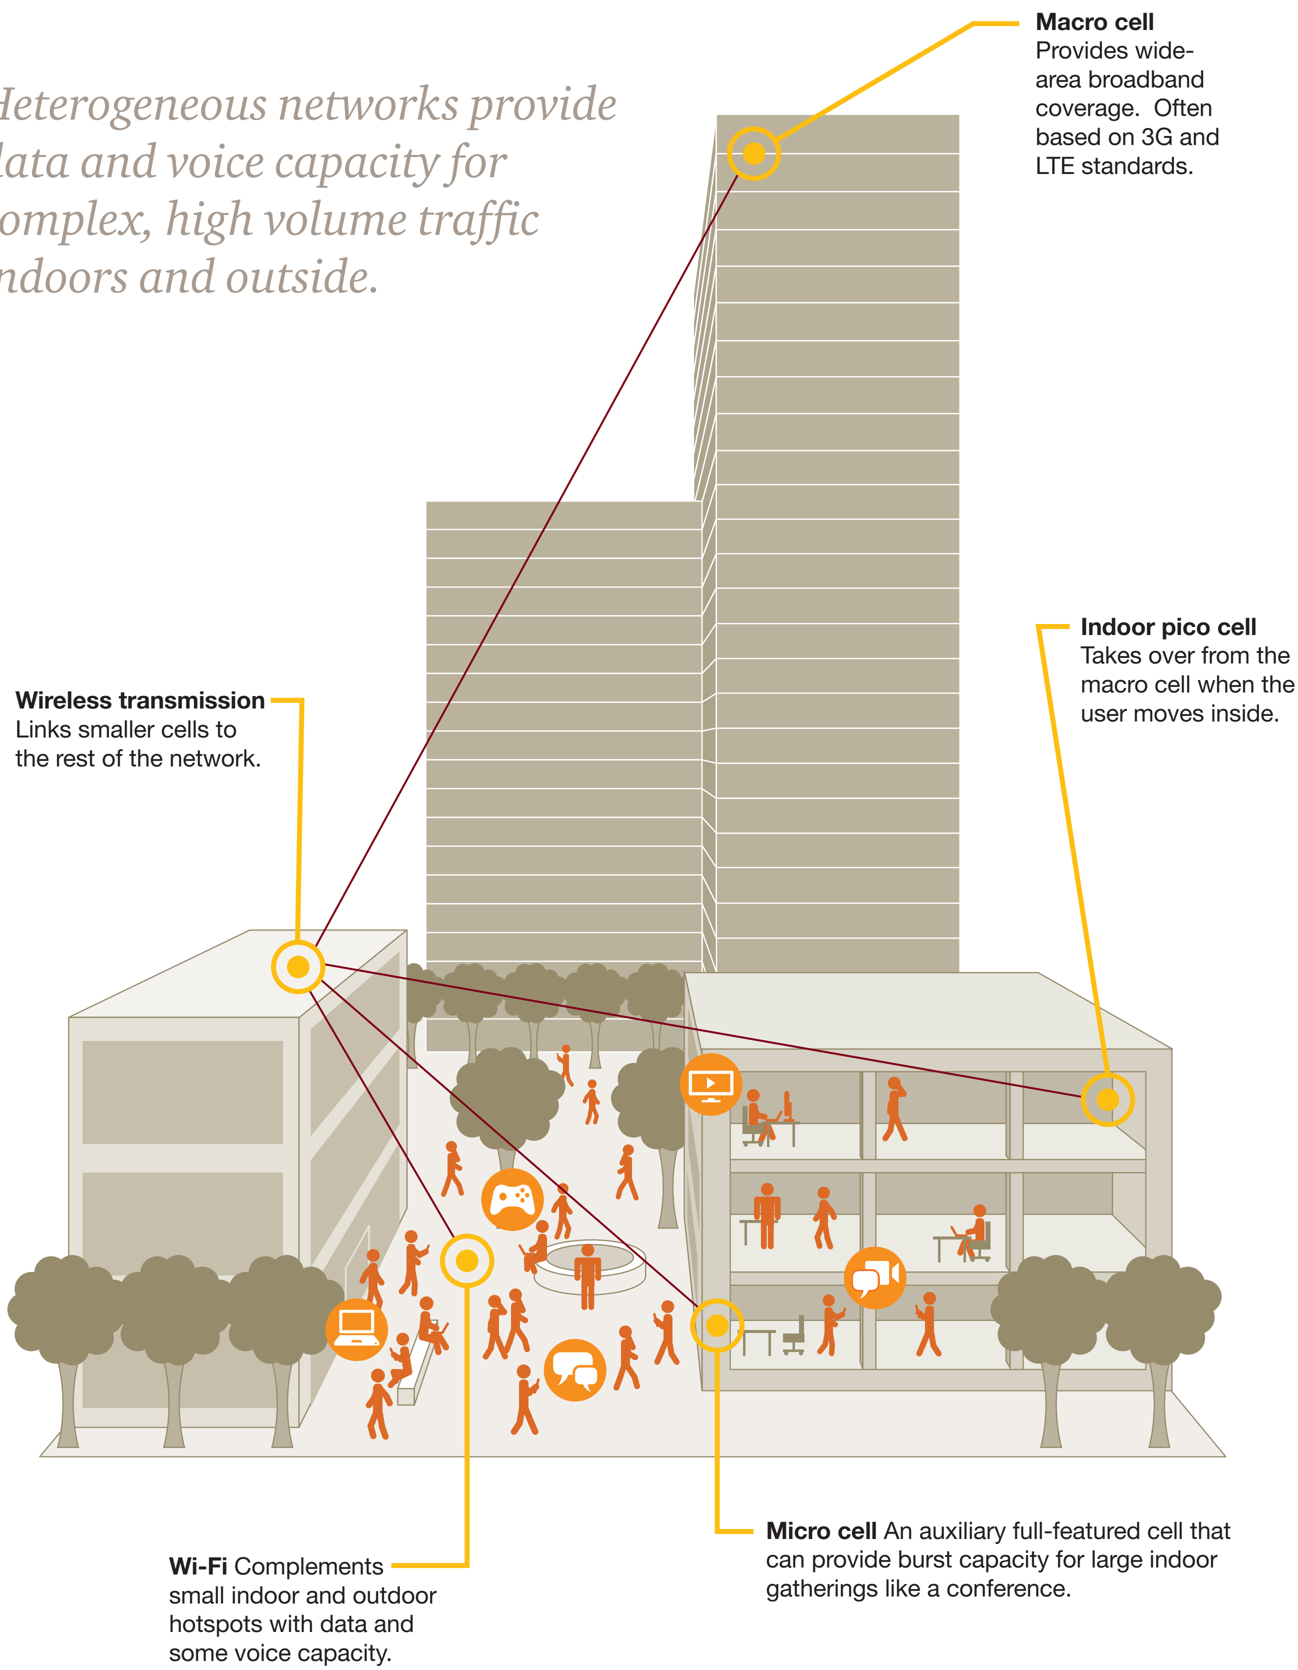

Source: PwC

Contextual services and capacity

Historically, network providers had a relatively simple idea of throughput. Throughput meant the average rate of successful message delivery through a given communications channel. More often than not, this meant the downlink side from the network to the mobile device. The volume of downloaded messages, the direction of message flow, and similar factors defined the networking industry, not simply in terms of the technology it adopted, but also the mentality with which it approached communications problems.

That model is changing rapidly as massive mobile data growth, a shift to visually oriented mobile data and the evolution to an app-centric usage model alter the mix of uplink and downlink mobile traffic patterns. Contextual services will combine and compound each force as end users generate large amounts of sensor data, captured media and applications that must upload rapidly to the network, process and then return almost immediately to the user as contextually relevant results.

At the most basic level, mobile communications networks are contending with unprecedented demand for capacity. Smartphones and connected devices pushed data past voice as the dominant source of mobile traffic in 2009. Ericsson reported in its June 2013 Mobility Report that total data traffic on mobile networks, in petabytes per month (1 petabyte = 1 million gigabytes), almost doubled between Q1 of 2012 and Q1 of 2013 (from just under 800 petabytes to just under 1,600 petabytes).[1] Cisco's Visual Networking Index also reported massive growth, calculating 885 petabytes per month of mobile data worldwide during 2012.[2]

These massive increases are happening as mobile video, image and visualisations are becoming the dominant mobile data types. According to Cisco, mobile video exceeded 50 percent of all mobile data traffic for the first time during 2012. As screens for mobile devices grow in size and resolution, even more visual traffic is expected to travel over mobile networks. According to Ericsson, mobile visual data is expected to grow by 55 percent annually until the end of 2019.[3] In 2013, the company reported that smartphone users who subscribed to music and video streaming services already consume more than 2GB of mobile data per month.

As important as the growth in mobile video volume, has been heightened user expectations for mobile video performance. For example, a 2012 study by Akamai Technologies and the University of Amherst on 23 million video streams from 6.7 million viewers showed that viewers start to abandon a video stream if it takes longer than two seconds to start, with each subsequent one second delay causing an additional 5.8 percent increase in abandonment. The study also showed that a viewer who experienced a failure in performance was 2.3 percent less likely to visit the same video destination during the same week compared to a user who did not experience a failure.[4] Consequently, superior performance for mobile video is no longer a special case for high-end users but is expected as a standard feature by the general population of mobile data users.

The shift in mobile devices from person-to-person communications tools to the primary interactive lens through which users connect to people, data, applications and services is driving mobile data growth. Regardless of the particular instantiation of a contextually aware communications session, the need for low latency, high throughput and blended data and communications all point to fundamentally different capacity models and mobile network architectures than those which brought cellular communications into mainstream life. ■

1,3 http://www.ericsson.com/mobility-report

2 http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html

4 http://www.akamai.com/dl/technical_publications/video_stream_quality_study.pdf

*Users **perceive network value** as a function of **application performance** as much as **voice quality**.*

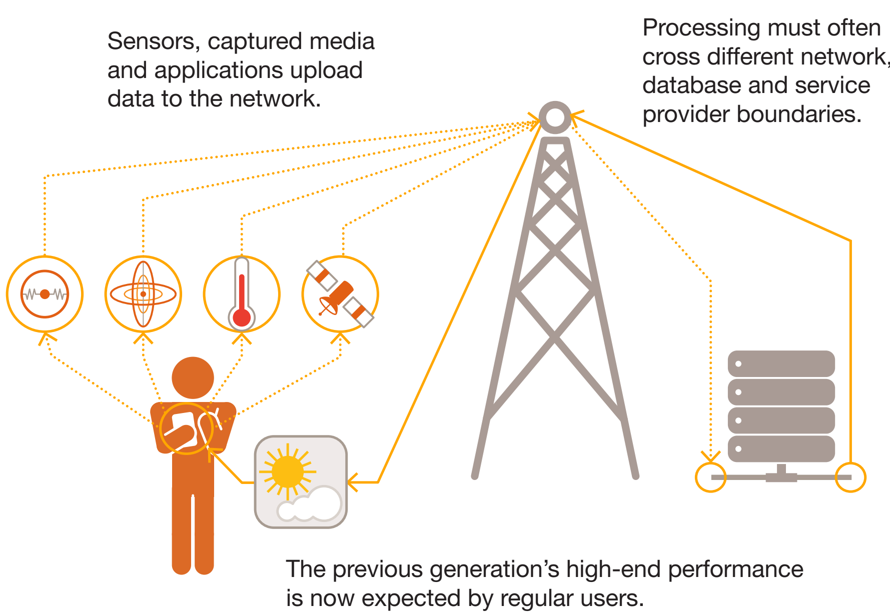

Source: PwC

Communications networks are increasingly the front door to the customer for third-party service providers that comprise a larger portion of modern economies.

differences in material and architecture inside structures that can affect radio wave propagation. Optical fibre moves captured radio signals between a central facility and the remote DAS antennas. This makes sense in densely trafficked areas such as an airport or a convention centre.

Whilst DAS is currently a mainstream indoor coverage strategy, it requires considerable expertise and investment by the CSPs to deploy and operate. Other distributed antenna strategies aim to turn mobile broadband antennas and their placement into a near plug-and-play proposition. This is the design philosophy behind Ericsson's Radio Dot System, which will launch commercially in late 2014. The actual active radio antennas or 'dots' weigh around 300 grams and deliver mobile broadband access to indoor users. Dots are connected and powered via standard Internet LAN cables that feed to floor-level radio units that all connect to a base station.

High capacity indoor coverage is not optional for CSPs that are faced with user expectations for instantaneous, reliable access to their apps and data wherever they are located. Equally important, the business case for LTE and small-cell radio coverage is being driven by the requirements of applications rather than voice. Along with providing more capacity, CSPs must provide smart capacity in order to prioritise different traffic streams and enable various business models. To make that happen, networks are becoming more programmable.

The network is a cloud

Communications networks are increasingly the front door to the customer for third-party service providers that comprise a larger portion of modern economies. At both the technical and business levels, CSPs will need to architect their networks to host ecosystems of third-party service providers engaging end users with various contextual experiences, at different price points and under different business models. To meet the rapidly expanding virtual context of users and the businesses that support them, CSPs are moving to more programmable networks.

The shift toward making networks programmable starts with the all-IP architecture of LTE. In contrast to the circuit-switched model of cellular communications that led to today, LTE supports only packet-switched services. The goal of LTE is to provide seamless IP connectivity between a client device and the data packet network without disrupting the user's applications whilst mobile.

Whilst LTE puts voice, data and applications onto a single delivery platform, new core network paradigms such as SDNs and its complement NFV are transforming how the network configures and operates on the inside. SDN and NFV use software to separate control of infrastructure elements from the underlying physical hardware to make a communications network operate more like a computing cloud. [See sidebar on page 7]

SDNs were pioneered in campus networking environments at the University of California at Berkeley and Stanford University in 2008. The purpose of SDNs is to allow administrators to shape communications traffic around different quality of service goals and/or business models without requiring admins to touch physical switches, routers or other hardware each time they want to make a change.

SDNs separate the part of network architecture that creates the map of nodes, links and addresses—the control plane—from the network architecture that makes decisions about what to do with inbound data packets (error correct, forward, reject, etc.)—also known as the data plane. This separation abstracts the physical hardware from applications and services riding on top of the network. Network administrators can make changes or add and drop features from a central location instead of having to hand code hundreds or thousands of individual pieces of equipment. In addition to abstracting the network, SDN architectures support a set of APIs that make it possible to implement common network services such as routing, security, access control, bandwidth management, traffic engineering, quality of service and other forms of policy management, any one of which can be custom tailored to meet business objectives within their own organisations or on behalf of other organisations.

Whilst SDNs emerged in campus and data centre networking environments, NFV had its origins amongst European CSPs that did not want to continue buying proprietary network appliances to run each new telecom service. Instead, CSPs wanted to launch so-called virtual network functions to run on virtual machines housed on standard servers. The European Telecommunications Standards Institute (ETSI) launched the NFV group to spur the development of interoperable products to address diverse use cases.

In practice, NFV decouples various network functions, such as network address translation, firewalling, intrusion detection, domain name service, caching, etc., from proprietary hardware appliances, so they can run in software. NFV is designed to consolidate and deliver the networking components needed to support a fully virtualised infrastructure—including virtual servers, storage and even other networks.

The table below compares some of the key points of SDN and NFV.

The immediate effect of all-IP infrastructures combined with network abstraction architectures via SDN and NFV is a communications network that allows approved applications to instruct network elements directly about their needs (routing, security, and performance, for example). Conversely, the network can broadcast its capabilities, state, analytics and other data to applications that want or need to access them.

The ultimate impact of this bi-directional information flow enabled by SDN and NFV is a more open and platform-oriented approach to networking. Communications networks become configurable services

Table 1: Key features of Software Defined Networks (SDN) and Network Function Virtualisation (NFV)

Category	SDN	NFV
Reason for being	Separation of control and data, centralisation of control and programmability of network	Relocation of network functions from dedicated appliances to generic servers
Target location	Campus, data centre/cloud	Service provider network
Target devices	Commodity servers and switches	Commodity servers and switches
Initial applications	Cloud orchestration and networking	Routers, firewalls, gateways, CDN, WAN accelerators, SLA assurance
New protocols	OpenFlow	None yet
Formalisation	Open Networking Forum	ETSI NFV Working Group

Source: http://www.sdncentral.com/technology/nfv-and-sdn-whats-the-difference/2013/03/]

that are accessed via APIs. Some embedded functions, such as firewalls, become customer controlled and tailored services in their own right. However, the same firewall service might also become part of a larger customer-facing bundle, such as a live video health counseling session, or a customer video conference with a tax professional.

The capability to provide à la carte network services based on application or business requirements rather than the specific configuration of physical infrastructure is called network slicing. In practice, a network slice is a logically isolated virtual network with its own service level guarantees that may span multiple data centres and network boundaries. Communications traffic within a given network slice is logically isolated from other traffic and can be further enhanced with firewall and encryption technologies. Fundamentally, the network slice concept brings to telecommunications the same model of on-demand, elastic resource allocation associated with cloud computing.

Figure 2: The network is the front door to the customer for many organisations

The network slice concept brings to telecommunications the same model of on-demand, elastic resource allocation associated with cloud computing.

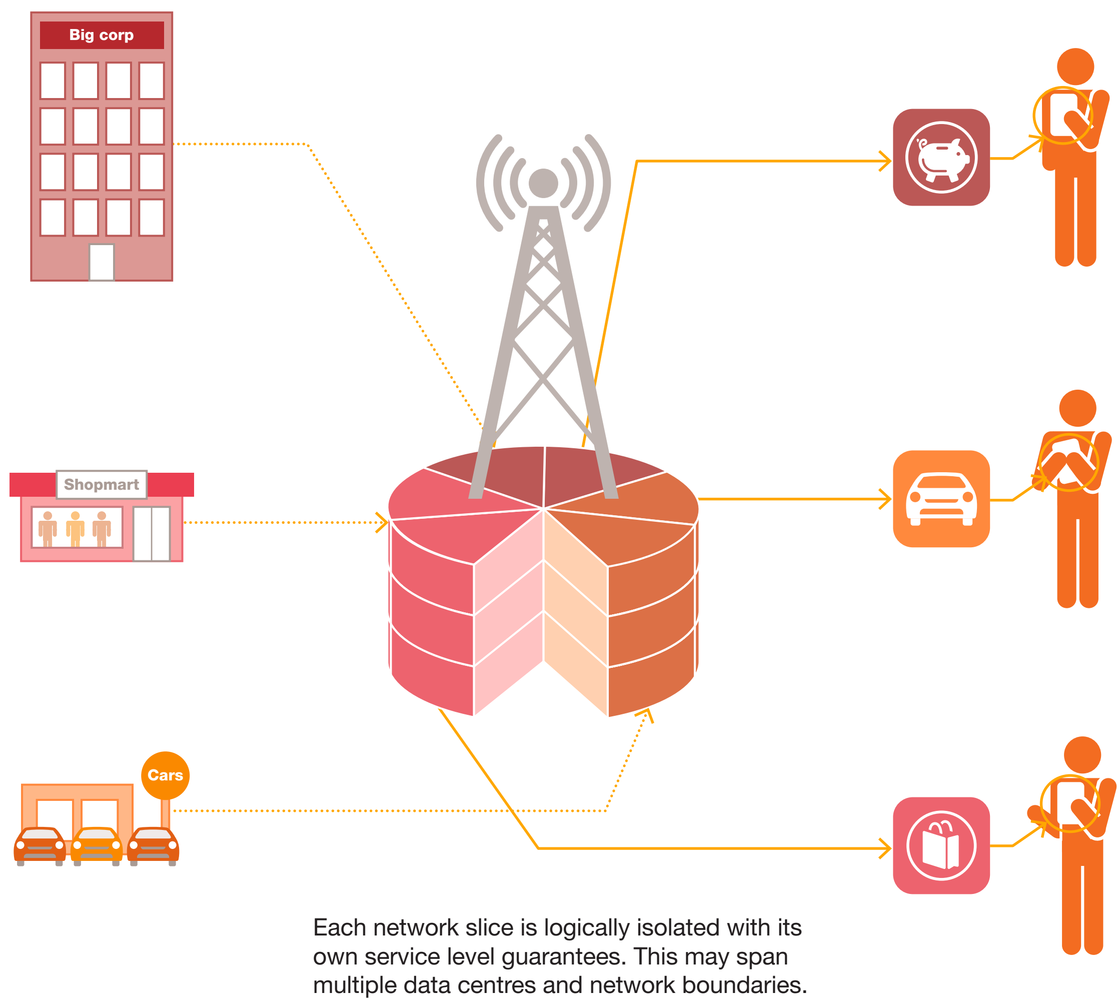

Source: PwC

SDN and NFV via OpenFlow and OpenDaylight

A number of open-source initiatives are driving the development of Software Defined Networking (SDN) and Network Function Virtualisation (NFV). Two of the most prominent are the Open Networking Foundation (ONF) [www.opennetworking.org] and OpenDaylight [www.opendaylight.org].

OpenFlow is an effort by the ONF to define a standard communications interface that separates the control and data planes of a networking architecture. As a result, the network appears to higher level applications and policy engines as a single, logical switch. With SDNs based on OpenFlow, enterprises and carriers gain vendor-independent control over the entire network from a single logical point, which greatly simplifies the network design and operation. SDN also greatly simplifies the network devices themselves, since they no longer need to understand and process thousands of protocol standards but merely accept instructions from the Openflow's SDN controllers.

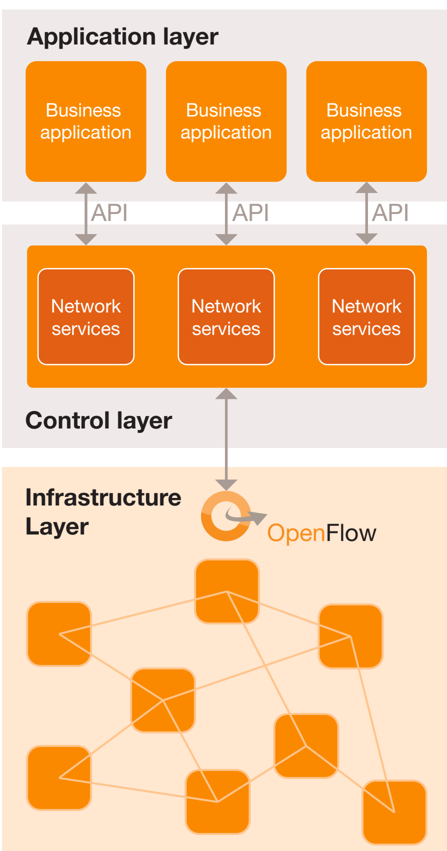

Source: https://www.opennetworking.org/sdn-resources/sdn-definition

Another important SDN standards player is OpenDaylight, which emerged from the Linux Foundation. OpenDaylight is an open-source project with a modular, pluggable, and flexible controller platform at its core. This controller is implemented strictly in software and is contained within its own Java Virtual Machine. As such, the controller can be deployed on any hardware and operating system platform that supports Java.

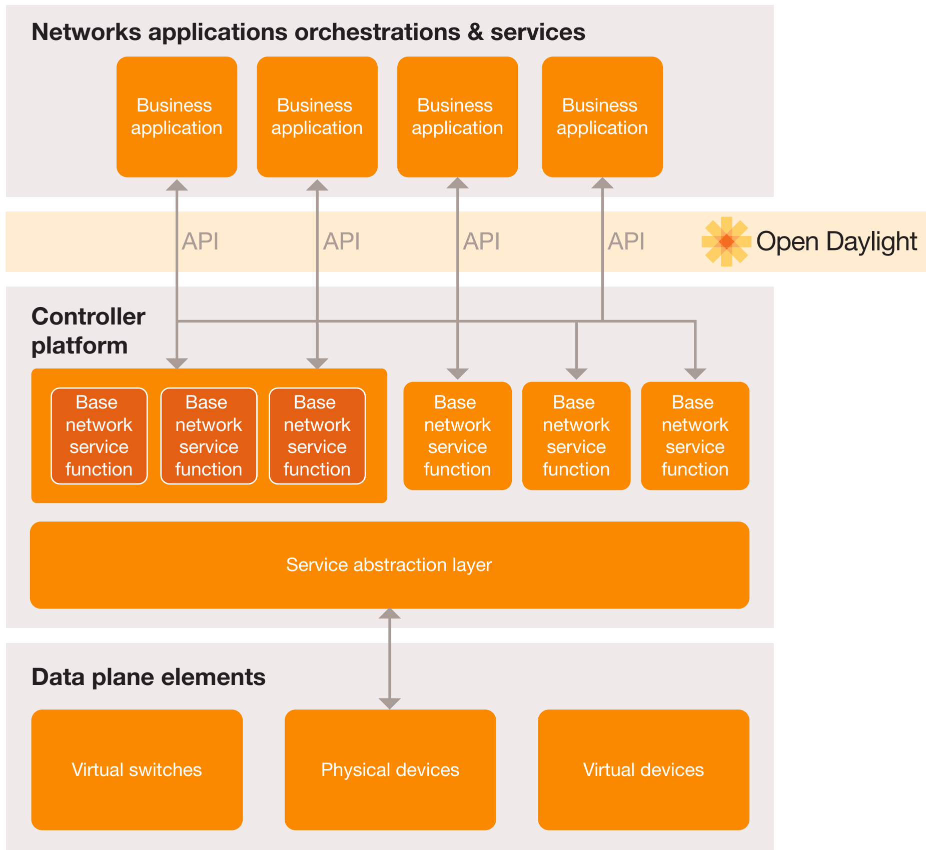

Source: www.opendaylight.org

Like OpenFlow, the OpenDaylight approach to SDN and NFV networking is a three-tier stack. The top level is called the Network Apps and Orchestration layer, which consists of business and network logic applications that control and monitor network behaviour. The middle layer, called the Controller Platform, is the framework in which the SDN abstractions can provide a set of common APIs to the application layer (commonly referred to as the northbound interface) whilst implementing one or more protocols for command and control of the physical hardware within the network (the southbound interface). At the bottom layer are the physical and virtual devices, switches, routers, etc., that make up the connective fabric between all endpoints within the network.

Officially, the two initiatives are not in direct competition. OpenDaylight will include support for the OpenFlow protocol, but will also be extensible to support other emerging SDN open standards, according to the OpenDaylight Foundation. The fact that many of the leading infrastructure and service provider organisations are founding members of both groups supports the case for complementary development. ■

At the business level, this means that CSPs can configure slices of the same network customised for different applications or industries. Application developers are able to access network services and capabilities without being constrained by the details of implementing them in physical infrastructure. In that sense, the game changer from being able to slice a network is the evolution of communications infrastructures from being just utilities to becoming true innovation platforms. According to Ulf Ewaldsson, Chief Technology Officer for Ericsson, network operators and third-party service providers will be able to access and direct networking resources under a cloud-based model similar to how they access storage and computation. "It means that developers can quickly build and deploy services outside of the operator's domain that are using the capabilities of networks from the device all the way to the data centre."

It's an app-driven world

The rise of virtual context means that CSPs must now assume that every communications device is simultaneously a computing and sensing endpoint for a user's personal cloud. That personal cloud, in turn, interacts with the surrounding physical environment, creates a user's proximity network and enables a portfolio of CSP and third-party clouds that deliver contextually relevant and intelligent experiences.

These contextually intelligent services will have significant uplink traffic as users generate massive amounts of data through device and environmental sensors, video and audio capabilities that were covered in the previous article. Aside from the low-latency required to provide the user with a quick and seamless experience, the network must route data traffic through the personal cloud of the user plus multiple clouds of various service providers depending on the user's current situation.

Consequently, mass-market expectations about network coverage and capacity for applications are matching the expectations that were previously restricted to high-end users. This suggests that the value proposition and competitive centre of gravity of communications networks will increasingly revolve around the needs of applications rather than voice or messaging, not to mention the value propositions and business models of third-party service providers.

Thus, virtual context is a user-driven rather than an infrastructure-driven technology landscape. The performance characteristics and value-creating activities of networks will evolve increasingly according to the requirements of users. In that sense, the expanding virtual context of users and the evolution of the networked cloud go hand-in-hand. In the next article in this series, we will explore the final layer to the contextual stack, the intelligence engines in the cloud that take data from a user's physical and virtual context and analyse it for predictive actions or suggestions.

Let's talk

If you have any questions about the Mobile Innovations Forecast or would like to discuss any of these topics further, please reach out to us.

Raman Chitkara
Global Technology Industry Leader
PricewaterhouseCoopers LLP
raman.chitkara@us.pwc.com

Pierre-Alain Sur
Global Communications Industry Leader
PricewaterhouseCoopers LLP
pierre-alain.sur@us.pwc.com

About PwC's Technology Institute

The Technology Institute is PwC's global research network that studies the business of technology and the technology of business with the purpose of creating thought leadership that offers both fact-based analysis and experience-based perspectives. Technology Institute insights and viewpoints originate from active collaboration between our professionals across the globe and their first-hand experiences working in and with the technology industry. For more information please contact Raman Chitkara, Global Technology Industry Leader.

About PwC

PwC helps organisations and individuals create the value they're looking for. We're a network of firms in 157 countries with more than 184,000 people who are committed to delivering quality in assurance, tax and advisory services. Tell us what matters to you and find out more by visiting us at **http://www.pwc.com**