# VIRTUAL NETWORKS

Telecoms operators have missed the platforms boat but hope to regain ground with network virtualisation. **RICHARD FEASEY** discusses the technology and regulatory implications of a powerful but potentially double-edged movement

In the traditional telecoms industry, technology and business strategy often occupy different worlds. Telecoms operators give their technologists and engineers remarkable freedom to pursue their dreams, only to find that the reality is different and less inviting. This was the case with the IP (internet protocol) revolution of the past 15 years, which left operators with lower costs and more flexible networks, but also without any control over the services environment, weaker relationships with their customers, and challenges in monetising their investments. Is the same thing about to happen again? I think it is.

This time the driver of change is not IP, but the network virtualisation movement or what Korea Telecom refers to as the 'transition to IT'. Once again, the technologists and engineers within the operators are enthusiasts.

Network virtualisation promises to reduce costs, both capital and operating, still further, to allow operators to run multiple logical or virtual networks over a single physical infrastructure, and to create opportunities for greater innovation and more competition among their suppliers. It does not require 5G to be realised, but key elements of whatever 5G vision you happen to subscribe to assume that network virtualisation will be adopted.

The logic is deceptively simple: split out the network intelligence from the underlying hardware, commoditise the latter and open source the former. In short, turn the traditional 'network' into a platform business. Traditional telecoms operators have mostly missed the boat on platform businesses so far, so perhaps this is their chance. If they succeed, operators would dump their traditional, vertical supply chain and replace it with two-sided ecosystem, with coders on one side and

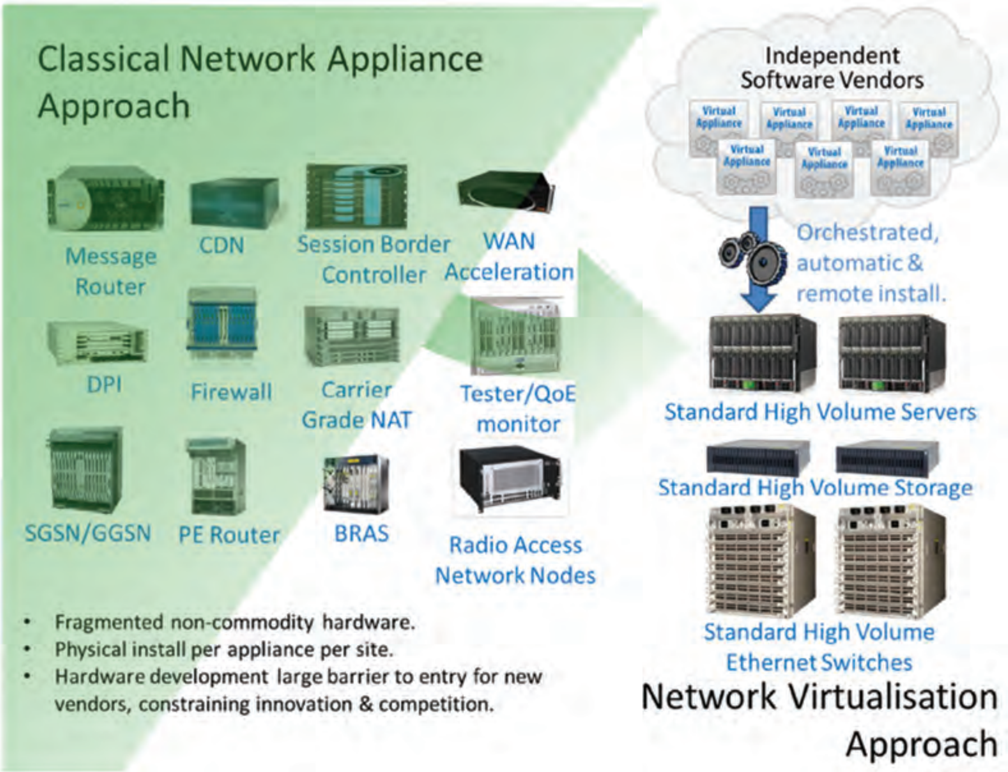


SOURCE: ETSI WHITE PAPER ON NFV (2012)

producers of standard IT hardware on the other. AT&T is clear that this is its aim:[1]

*"AT&T expects to develop key software resources in a way that they can be openly used, and cannot be lost through the acquisition or insolvency of a vendor partner. This pivot will enable AT&T to do business with start-ups and small businesses that we might have deemed too risky in the past. While they may not always endure, small businesses demonstrate the large fraction of innovation and agile development in the marketplace and enabling the company to do business better with these small companies is a key element of Domain 2.0."*

This is a transition in which the software that orchestrates the functions performed by what we refer to today as 'the network' and the underlying physical hardware of that network are decoupled into two sides of a platform, just as services and networks were decoupled with IP. The underlying architectural principles are generally known as software defined networking (SDN) and network functions virtualisation (NFV).

## TECHNOLOGIES EXPLAINED

SDN refers to the decoupling of the control plane from the packet forwarding functions (or 'data plane'). These functions are today integrated within the same network components, such as routers, in which communication between the control and data planes is undertaken using propriety code written by the vendor. With SDN, the aim is to use standardised hardware that could be supplied by many different vendors that is then controlled by software, which anyone could write, through a standardised interface.

Just as IP enables us to combine traffic from different sources into a single multiplex rather than running a series of discrete networks, so SDN would enable a wide range of logical 'networks' over common physical resources. The common physical resources of the network, including the radio spectrum, could then be dynamically allocated and reallocated between the different network slices without any changes to the underlying hardware.

This capacity to create different 'slices' over a common network substrate is a key feature of the 5G vision and would allow internet of things (IoT) applications that need only narrowband, low power, delay tolerant conditions to run alongside gigabit wireless fibre applications over the same physical network.

NFV envisages virtualising network equipment as software so as to run a wide variety of applications on virtual machines hosted on standardised IT hardware. Many functions that are currently performed by dedicated components in the traditional network – firewalls, load balancing, intrusion detection – could be virtualised. In each case the idea is that the proprietary hardware is replaced by a 'virtual machine' which is hosted on a generic IT or cloud infrastructure. SDN can be deployed without NFV, and vice versa, but most operators envisage deploying both and the concepts appear highly complementary.[2]

There are many benefits from SDN and NFV, most of which are common to both. Shared network assets, including spectrum, can be managed more holistically and efficiently. Routing decisions that are taken by an omnipotent network controller will ensure more optimal utilisation of the network as a whole, as well as allowing for different service policies to be applied for different flows or sequences of packets. This architecture also allows for cloud computing and IT capabilities to be hosted at the very edge of the network – a concept often known as 'mobile edge computing' – which is a prerequisite for some of the very low latency applications envisaged for 5G.

> **"Shared network assets, including spectrum, can be managed more efficiently."**

Orchestrating networks in this way improves efficiency and lowers costs. SDNs ought to be easier and faster to upgrade in order to support new services, since there is less dependency on hardware that is difficult, expensive and slow to replace, and more dependency on software, which is easy, cheap(er) and quick. When hardware does need to be replaced or augmented, standard IT hardware and open interfaces ought to increase competition and reduce dependency on particular vendors, driving costs down: Google claims 95% utilisation on its 'virtualised' backbone (compared with industry norms of 30-50%), while AT&T has announced plans to 'virtualise' 75% of its network by 2020 and predicts that it will reduce hardware costs by 30-40% and operating costs by 50%.

Operators such as AT&T, China Mobile and Korea Telecom are announcing ambitious plans to virtualise their networks. But challenges and uncertainties remain. There is a huge amount of proprietary network equipment to be replaced or to be integrated into the new architecture. Operators need staff to develop software, but many traditional telecoms engineers lack these skills and those who have them are in short supply. There are security concerns and other technical risks of the kind associated with any massive IT project.

## WHO WILL BENEFIT?

But the question is surely not if but when the virtualisation of networks will happen. The virtualisation movement has powerful backers, and we have seen that the potential benefits are huge. The more interesting question is who will reap the benefits when it does. Of course, the telecoms operators assume that they will sit at the heart of the platform.

One clue might be provided by looking at who is driving the technical work required to make virtualisation happen. Existing platform businesses such as Facebook and Google are the key movers behind various industry initiatives which are seeking to reduce hardware costs (which means reducing energy consumption in data centres)[3] but also efforts to virtualise functions currently performed by switches in telecoms networks so that they can instead be hosted in data centres,[4] as well as promoting the standardisation of software interfaces between applications and network operating systems and between these systems and hardware.[5] These firms have a long-standing interest in driving down the costs of internet access, which network virtualisation promises to do. But they also have skills in coding and in building platform businesses, and it may be these skills which matter most in determining who controls the platform in future.

All this needs to be seen in the context of other trends that are changing our view of what it means to operate a telecoms network or to be a telecoms operator in the future:

● **Infrastructure outsourcing.** Mobile operators have been divesting towers to tower companies ('the towerco') for many years. In the US, 80% of mobile towers are already owned by an independent towerco rather than by the operators. In Europe, Middle East and Africa it is closer to 20%, but the trend is in the same direction. Operators are also outsourcing significant parts of their active networks (to traditional vendors such as Ericsson and Nokia) in order, they hope, to benefit from economies of scale, reduce the operational complexity of their businesses, and perhaps for financial engineering reasons as well. Whether virtualisation will exacerbate this trend or reverse it is unclear to me. There is a clear potential for tension between the efforts of traditional vendors to extend their control over network assets, and the aims of the virtualisation movement and of operators like AT&T, seeking to reduce their dependency on such vendors.

● **Sales channels.** Trends in retailing and distribution in the telecoms industry are changing too. Many operators have long relied on third parties for retail distribution, particularly in mobile and particularly for prepay services. On the other hand, there has also been a trend in some markets towards reintegration, with network operators buying out third party distributors to obtain more control over their sales channels. Some have long speculated about whether online distribution of mobile devices will replace traditional retailing, as it has for other consumer electronics, but the long queues outside Apple stores suggest otherwise. We are some way off that today.

> **"Virtualisation will happen – the more interesting question is who will reap the benefits when it does."**

● **Spectrum.** Ownership of radio spectrum has traditionally been a core function and exclusive asset for mobile operators. But this appears likely to change as software enables the sharing of spectrum. 5G envisages the availability of significant tranches of unlicensed or shared spectrum: in the FCC's Spectrum Frontiers proposal for millimetric spectrum (above 24 GHz), almost twice as much spectrum (7 GHz) is being made available for unlicensed purposes as for traditional licenced applications.

We also know that the transition to IP meant an increase in services competition, the globalisation of the market and the disintermediation of the network operators. If those trends are also a feature of the 'transition to IT' then the 'network' will become increasingly fragmented, ownership and control dispersed, and boundaries redefined. AT&T accepts that some element of control might pass from the operators to customers and to third parties, although it still assumes that it (AT&T) will retain control of the 'platform':

*"NFV can readily be applied to the control and management plane in addition to the data plane. This allows virtual networks to be created and managed by end users and third parties using the tools and capabilities heretofore reserved only for native network operators. Domain 2.0 comprises more than simply a network or service architecture. It requires appropriate business practices, a supplier and software ecosystem, a software-savvy planning and operations organisation, and management willing to try alternatives and fail fast."*

Some of today's telecoms operators – those with the resources of AT&T or China Mobile – may be better placed to beat the software giants at their own game than they were 20 years ago when the IP transition began. But it is difficult to avoid the suspicion that virtualisation plays to the strengths of existing platform businesses rather than to those of the telecoms operators. The key question may be whether the control of the software sits better with the owner of hardware or with the provider of the applications (or the user of the applications).

If the project to turn central offices into data centres[4] is taken literally, then large corporates and those involved in IoT may literally buy a 'network as a service' from a data centre provider. This is a market where the internet and IT companies such as Amazon, Microsoft and Google, not the telecoms operators, dominate.

If this is right, there are several lessons for policymakers and regulators to be drawn from the history of the IP transition which are likely to be relevant to the IT transition and network virtualisation:

- Policymakers will want to ensure that incumbent players, both traditional telecoms equipment vendors and traditional network operators, do not try to erect barriers to virtualisation. Such barriers might involve insisting on proprietary code and standards rather than open interfaces and open source software. Bundling hardware and software, or hardware and hardware, or software and software together, could also be barriers. The standards bodies for NFV and SDN – most of which sit outside the traditional telecoms groups such as the European Telecommunications Standards Institute (ETSI) – could play an important role in avoiding this.
- Although the incumbent players may attempt to move from hardware into software, the internet companies may also attempt to gain from moving from applications into network software. Application providers might, for example, design their services so they would only run properly over their own virtual network and not over those of rivals. This would be a complex area to police and would necessitate a serious re-examination of existing net neutrality assumptions to recognise that the decoupling process means that today's internet access provider may no longer control the network in future. (These are the terms used in today's European net neutrality rules. I assume net neutrality rules in the US may be revised for other reasons following the election of President Donald Trump.)
- The virtualisation of networks and decoupling of software from hardware may also lead regulators to worry much less, or not at all, about competition in hardware. The imperatives – better asset utilisation and cost reduction – that are already driving the formation of the towerco, network consolidation and spectrum sharing, might end with a reversion to a single physical network over which competing virtual networks are then deployed. In a decoupled world, duplication of physical network assets would yield few benefits and lots of costs, while consolidation of these assets should not mean any loss of competition in the virtual realm. Networks could then be hosted in enormous data centres, with their location determined by energy costs or the regulatory environment, rather than the location of the customers they serve or the physical assets or hardware over which they run.
- Regulators will face the same issues of supervising and regulating networks as they now have with regulating IP services that are provided by firms which operate remotely and on a global basis. They have already lost a significant degree of control over IP services, but would now face the prospect of losing control over the virtual networks as well (including being unable to shut them down). Being able to regulate the towers and hardware that remain within a country will not count for much if control of the network sits elsewhere and there is no exclusive rights to spectrum which a government could withdraw.

> **"Policymakers may need to think more about how the value chain works in the interests of users.**

### EVALUATING THE VALUE CHAIN

The history of the IP transition also suggests that there are some other significant risks ahead. One will be the challenges which the participants in a long and complex value chain face in coordinating their interests and activities in order to produce what economists call complementary goods. The classic example with IP services has been the ongoing tensions between the telecoms network operators and IP service providers about the costs (in terms of additional network capacity) which those services impose on the networks and how they might best be recovered.

So far, the value chain as a whole has been able to deploy more network capacity and broadly to meet the demands that new applications have presented without the internet collapsing. However, we now seem to be moving into an ever more complex world in which the demands of applications on 'networks' will be much more heterogeneous and the degree of coordination between different entities required to satisfy them will be much more challenging.

Policymakers have so far been able to largely ignore the complaints of the different participants in the value chain, and they can probably continue to do so while the overall system continues to deliver. My sense is that they watch nervously as Comcast and Netflix argue about peering arrangements or European operators threaten to block adverts, knowing that if something really did go wrong and services were badly disrupted, then they probably have few effective tools with which to fix it. Network virtualisation might require policymakers to think more seriously in future about how to ensure the value chain works effectively in the interests of users.

Whether Google, Facebook or Amazon end up running vast global networks instead of today's telecoms operators I do not know. If they do, the current enthusiasm for network virtualisation among the engineers inside the traditional telecoms operators (as well as their traditional equipment vendors) may turn out to produce an unwelcome surprise for their colleagues in strategy, as well as a new set of challenges for policymakers.

*__RICHARD FEASEY__ is an independent consultant and an associate at Frontier Economics, having previously worked for a variety of global telecoms operators over a period of 25 years. The views expressed in this article are personal and should not be attributed to anyone other than the author himself.*

**REFERENCES** **1** AT&T Domain 2.0 Vision white paper. soc.att.com/2gZSu7T **2** There are numerous introductions to SDN and NFV from vendors and industry groups. The ETSI 2012 white paper on NFV is often cited – see bit.ly/1tsRH2A – and I have found the AT&T white paper (ref 1) particularly useful. **3** The Open Compute Project is focusing on reducing data centre costs – opencompute.org; see Open Cellular on cellular network costs – bit.ly/29ioWvx; and the Telecom Infra Project on core network, backhaul and access – telecominfraproject.com. Facebook is active in all three groups. **4** The aim of the Central Office Rearchitected as a Data Centre (CORD) project is self-explanatory and has both Facebook and Google (as well as AT&T and Verizon) as participants. opencord.org **5** OpenDaylight claims to be the largest open source SDN industry player. opendaylight.org